

Library of Congress

Experimental Access to Digital Collections and Data

Final report: project overview, findings and
recommendations

Version 1.0
March 2022

Contents

Summary	6
Discovery and Ideation	9
Project goals	9
Target audience	9
Personas	11
Educators	11
Activist/community leaders	11
Data journalists	12
Undergraduate creative / art students	12
How Might We?	12
User Insight 1: Search and Ask (a librarian)	13
Summary	13
How Might We	13
User Insight 2: Collaboration and Impact	14
Summary	14
How Might We	14
User Insight 3: Trust	15
Summary	15
How Might We	15
User Insight 4: Today Problem	16
Summary	16
How Might We	16
Ideation and prototype generation	17
Technical and practical constraints	18
Infrastructure and time-to-prototype	18
Agile prototyping	18
No humans in the loop	19
Prototype 1: Fat Item Pages	20
Problems and proposed solution	20
User Insight 3: Search and Ask a Librarian	20
User Insight 4: Today Problem	20
User Insight 2: Trust	20
Proposal	20
Process: Wireframing, design and further ideation	21
User Stories	22

Design/UX features	24
Popups/previews	24
Interactive facets for similar objects	25
Generous interfaces / embedded links	25
Embedded IIIF based viewer	26
Use of maps and/or timelines	26
Development	27
Iteration 1: "Skinning the LoC JSON API"	28
Iteration 2: "The Fat Item Page"	31
Popups and previews	32
Entity information	34
Technical approach	35
User insights	36
Information-dense pages are welcome	36
People, locations: 'incredible useful' for educators and students	36
Provenance matters to people	36
Citation	36
Usability notes	37
Main image	37
People and location	37
Popup preview	37
Find similar items	38
Navigation and layout	38
Rights and access	38
Prototype 2: Aggregation and Clustering	39
Problems and proposed solution	39
User Insight 3: Search and Ask a Librarian	39
User Insight 4: Today Problem	39
Proposal	39
Process: Wireframing, design and further ideation	40
Design/UX features	42
Visualization of metadata or entity driven facets	42
Visual navigation	42
Summary information	42
Development process	43
Code Sandbox Examples	44
Approaches we ruled out	44
Dimensionality reduction and K-Means clustering	44
Facet co-occurrence matrix	44

Approaches we kept in	46
User insights	47
Search behaviors	48
Personalization	48
Curation	48
Clicking on an item is a commitment	48
Labels for all	49
Search vs. discovery	49
Usability notes	50
Interactive visualizations	50
Search results list	51
Filters	51
Enrichment and APIs	52
Approaches	52
Named entity recognition	52
Entity linking	52
Entity enrichment	53
Technology stack and queue based processing	54
Order of operations	54
Signals and tasks	54
Asynchronous task processing	54
Search	55
Enrichment and APIs: Risks and challenges	56
LoC APIs	56
Inaccurate entity linking	56
“Looping” and idempotency	57
Prototype 3: Tagging and Sharing	58
Problems and proposed solution	58
User Insight 1: Collaboration + Impact	58
Proposal	58
Process: Wireframing	59
Tagging	59
Sharing	60
User insights	61
User-generated tags upfront	61
A complement to fat item pages	61
Quick feedback	62
Great for learning concepts from users	62
Trust and value	62

Managing volume	62
Why stop at tagging?	62
Rewarding users	62
Challenges and lessons learned	63
Data Retrieval and Enrichment	63
What went well	63
Challenges	63
Data quality, accuracy and provenance	64
What went well	64
Challenges	64
Accuracy of the data	64
Data provenance	64
Contested, sensitive or pejorative terms	64
User Research	65
Challenges	65
Discovery and Ideation	66
What went well	66
Challenges	66
Development approach	66
What went well	66
Challenges	66
Lessons Learned	67
Accuracy, provenance and sensitive themes	67
Performance and scalability	68
Usability of the prototypes	68
Conclusions and potential next steps	70
Conclusions	70
Potential next steps	71
1. Enrichment	71
2. Performance	71
3. Visualization	71
4. Personalization	72
5. Discovery vs. Search	72
6. Research wider audiences and continue learning	72
Appendix 1: Prototype “pitches”	73
Cluster and aggregate search	73
For	73
Their problem(s)	73

User Insight 3: Search and Ask a Librarian	73
User Insight 4: Today Problem	73
Proposal	74
How can it help the Library of Congress?	75
How might it work?	75
Resources and supplementary material	76
Visualization	76
Other issues and considerations	76
Fat Item Pages	77
For	77
Their problem	77
User Insight 3: Search and Ask a Librarian	77
User Insight 4: Today Problem	77
User Insight 2: Trust	77
Proposal	77
How can it help the Library of Congress?	78
How might it work?	79
Tagging and sharing	80
Who is it for	80
Their problem(s)	80
User Insight 1: Collaboration + Impact	80
Proposal	80
How can it help the Library of Congress?	81
How might it work?	81
Tagging	81
Sharing	81
Resources and supplementary material	82
Appendix 2: User research and usability test methodology	83
Sessions	83
Participants	83
Approach	83
Test script	84

Summary

The Library's initial vision for this project was to

- Identify the needs, and especially the *unmet needs*, of library users with respect to the library's digital collections
- Assess, with respect to these user needs, the experimental digital collections interfaces produced by the LC Labs and other groups, and
- To evaluate, through user research and new prototypes, new and innovative approaches to how the library might address these needs

This report primarily covers the design and development of the experimental prototypes, and the corresponding user research. It also provides some recommendations for future avenues for exploration, both tactical and strategic, to provide better access to the library's digital collections.

Our approach for this phase of the project spanned 6 stages:

- [User research](#) and the development of user insights
- [How Might We](#) sessions to identify opportunities
- [Ideation and Sketching](#) sessions for initial prototype suggestions
- Prototype selection with LoC stakeholders
- Development of prototypes and iteration in response to feedback
- User testing and user validation

From the Ideation, sketching and prototype selection processes we shortlisted three prototype concepts:

- ["Fat" item page](#): looking at presenting a rich view of individual items
- [Aggregation and clustering](#): looking at clustering and visualization multiple items in aggregate
- [Tagging and sharing](#): looking at user created data and sharing

These prototypes were developed over a period of approximately 2 to 3 months and, after feedback and suggestions for improvements by LoC staff, were taken forward for user testing and evaluation.

In terms of user validation and feedback, our approach looked to get both feedback on the actual designs and to learn more about the users' mental model through deeper questioning.

We conducted 6 one-hour interview sessions to walk participants through the three prototype ideas we developed. All three concepts were found useful to the participants to varying degrees. Of the three, the fat item page generated the most interest and engagement. Overall participants were excited to see LoC working in these interconnected areas.

We have suggestions [below](#) about how improvements might be made and further opportunities explored, building on similar concepts in future, including approaches to improving performance and scalability, visualization, data transformation and enrichment.

Discovery and Ideation

Project goals

Early on in the discovery phase we established three goals with the LoC Labs for this project:

1. Provide evidence-based recommendations for making collections more available and usable
2. Understand how to reach users we at LoC are currently not serving
3. Inspire and inform out LoC colleagues throughout the Library

Target audience

These project goals then informed which audiences we needed to investigate through focused user research activities. Making use of the [inverted pyramid of user engagement model](#) developed by The National Archives in the UK, we looked to identify which overlaps existed between audiences, the project goals and The User Experience Continuum as described in [LoC's digital strategy](#): the project goals aligned well with the LoC's Discovery state of that continuum:

Our users should be able to find what they are looking for easily and as intuitively as possible. Current user experiences—shaped via search engines, social media platforms, available metadata, and always-connected devices—have led to an expectation that relevant information will rise to the top and in many cases present itself to a user even without asking. While this is a challenging expectation for any organization to meet, in some ways this is where libraries excel. Because of our expert staff, we are well-suited to expose interesting relationships, enable "serendipitous discovery," and facilitate creativity.

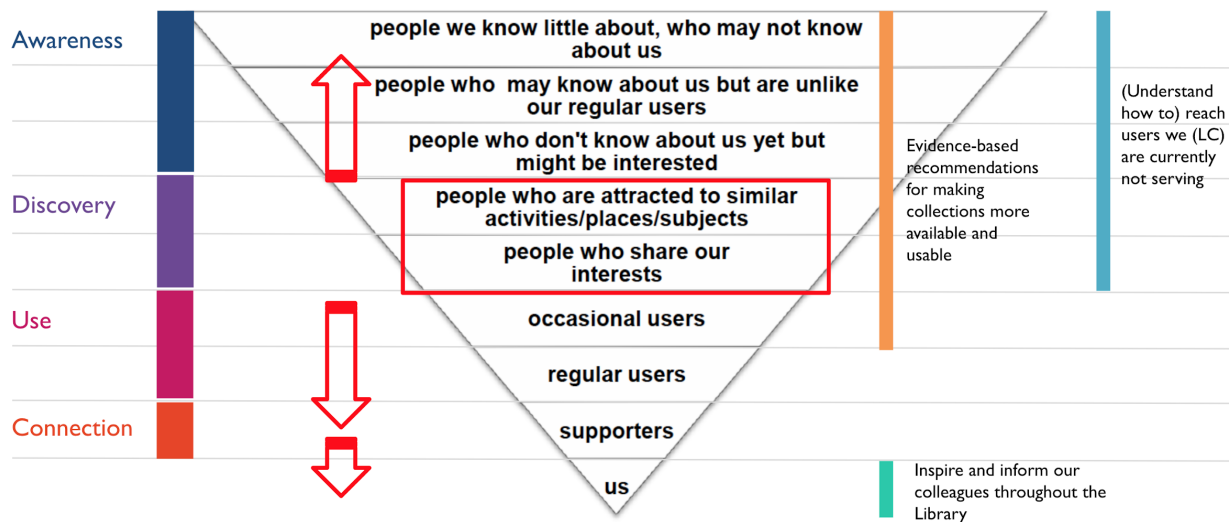


Figure. User engagement model.

The resulting alignment pointed us to focus our initial user research efforts to:

Users who might develop applications, share, learn, teach, connect or otherwise communicate with Library-like materials who are not currently connected to the Library.

After a workshop with LoC to explore a few user groups that could fall under that category, LoC's product owners decided to focus the research on four personas/user types:

1. Educators
2. Activists/community leaders
3. Data journalists
4. Undergraduate/ creative students

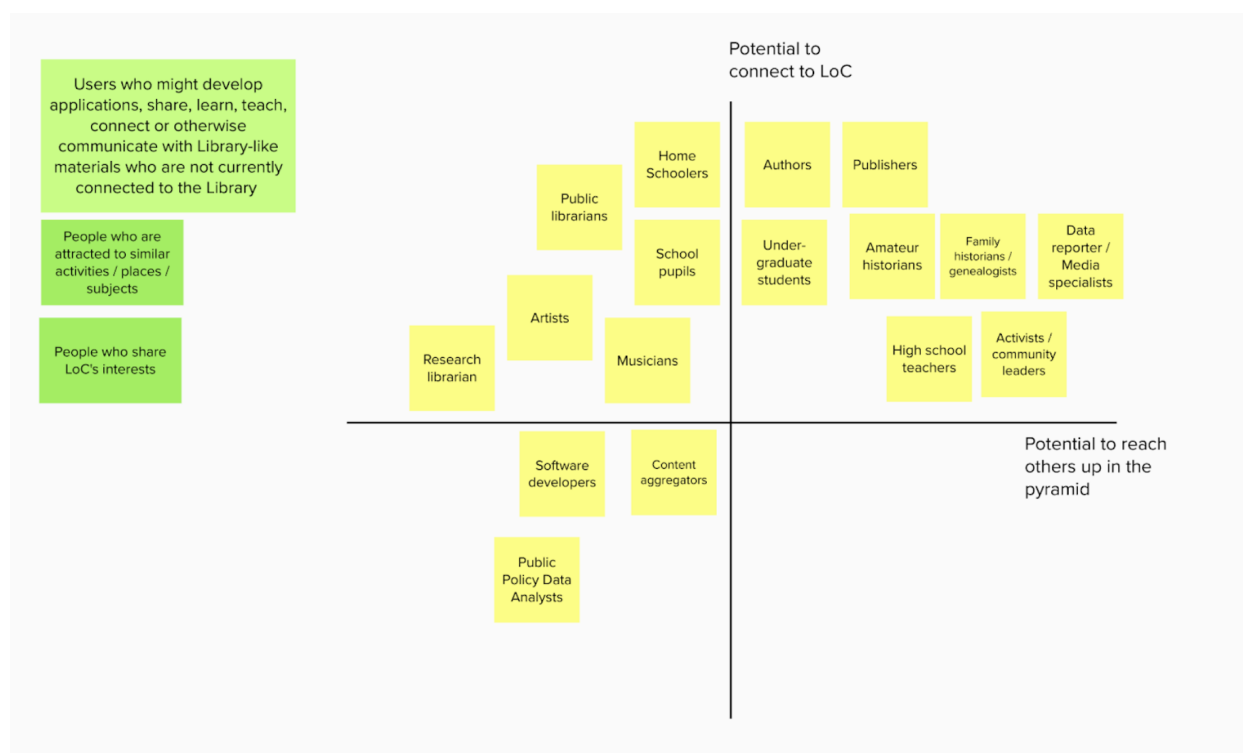


Figure. User groups explored during discovery.

Personas

Educators

They might use library-like resources to find alternative, more creative ways for their students to learn. They might expand opportunities for students to develop applied skills through logic/computational thinking. Primary and high-school educators were mentioned as well as formal and informal ones. Professional K-12 teachers that might use library-like resources to find alternative, more creative ways for their students to learn, for example, using these resources "as data" to help teach data skills or data literacy.

Activist/community leaders

They might use library-like resources to inform awareness of a community, history, or relationships, and in the process, change a community, aiming to make it a better place.

Data journalists

They might use library-like resources and data to produce attractive, informative visualizations with cutting-edge techniques to accompany data-driven stories.

Undergraduate creative / art students

They might use library-like resources as inspirational sources and, through their creative work, highlight the potential that the library holds as a complex digital space of knowledge and cultural development.

With these four personas agreed, we conducted an initial round of generative research to develop a deeper understanding of these users in order to find opportunities for solutions and innovation.

See [LoC Labs: User Research Report \(slide deck\)](#) for a summary of the findings.

How Might We?

The Discovery process consolidated a number of *How Might We?* (HMW) statements around a number of key User Insights drawn from user research.

1. Search and Ask (a librarian)
2. Collaboration and Impact
3. Trust
4. Today problem

By turning problems into opportunities, *How Might We?* statements help clarify particular areas for exploration and further research and development.

The key user insights and the *How Might We?* statements generated from each insight during a workshop with the LoC Labs team were as follows.

User Insight 1: Search and Ask (a librarian)

Summary

Often, when users have questions they want LoC to answer, they start by using the online search functionality but, when they fail to find what they are after and get stuck, they check with the Ask a Librarian team. Users struggle to make sense of search results, and lack an understanding of what the Library holds and what it doesn't. The material can present barriers too: it can't be accessed remotely as it might be under copyright (you need to be at the Library), and collections are siloed and require searches in multiple places. All of this can feel frustrating for users. Instead, ideally our user can determine quickly whether the Library might have answers to their questions and rely on Ask a Librarian only when all automated paths such as search have been extensively and clearly explored.

How Might We

- Search User Experience
 - How might we provide clear paths to the refinement of search?
 - How might we improve the relevance of search results?
 - How might we make search a more intuitive experience?
 - How do we remember what it is like to be a beginner?
- Search Paradigm:
 - How might we extend search beyond the catalog record paradigm?
 - How might we get rid of search result pages altogether?
- Empower the user:
 - How might we demonstrate what is often used in relation to the user's question?
 - How might we encourage users to engage with existing support resources?

User Insight 2: Collaboration and Impact

Summary

Currently, users recognise the need to work with individuals inside and outside their organization to create impactful work. Ideally our user can collaboratively and efficiently work with others, engaging with dense topics, being empowered to learn along the way, sharing their contributions, receiving feedback from other experts and being able to see the positive impact their work has on other people. All that keeps them motivated and engaged.

How Might We

- Supporting Connections
 - How might we connect people with shared interests?
 - How might we connect non-expert users with expert users?
 - How might we connect users to answer each others' questions?
 - How might we enable users to support each other?
- Tooling and facilitation
 - How might we connect groups in technical and physical spaces?
 - How might we provide a digital space in which collaboration can happen?

User Insight 3: Trust

Summary

Our user wants trustworthy data and information. Currently, they rely on their own expertise when faced with conflicting information. Finding trustworthy information is difficult for our user. Trusting what they find, unless they have direct experience with it, is also a challenge. Ideally our user can rely on organizations such as LoC to help make their information-driven decisions and arguments. By being supportive of others, LoC helps build trust which leads to receiving support in return.

How Might We

- Trust / Information Ecosystem
 - How might we enable collaboration with other organizations both in the private and public sector to raise the LoC profile as a source of trustable information?
 - How might we facilitate comparison to information in other collections or other organizations
 - How might we bring expert opinion and scholarship in combination with resources to build trust in the quality of that information?
 - How might we show how a resource has been used for other publications?

User Insight 4: Today Problem

Summary

Currently our user has a sense of urgency when dealing with whatever information needs they have at hand. Today problems get in the way of being able to give the time and attention necessary to create anything that isn't absolutely necessary for meeting immediate challenges. While our user has access to the resources on the LoC site 24/7, they do not think about it or come across it when looking for data and information. Collections are not optimally organized and still require a fair amount of work to frame or reframe in the needed context. Ideally our user can quickly and effectively access content at LoC in a way they can grab it and move on to the next problem of the day.

How Might We

- LoC as Primary Resource
 - How might we make sure users are aware when new info becomes available?
 - How might we better integrate LoC resources in broader web search (knowledge graph, etc)
 - How might we link the use of Congress.gov & CRS reports (very up to date info and use) to other resources the Library has.
 - How might we make LoC a first-stop when looking for information in a hurry?
 - How can research and LOC search be woven into a users daily internet usage?
 - How to use brief collection stories to get beginners interested in primary resources?

Ideation and prototype generation

These core insights into the problems users face and the opportunities that might be provided were the key input into an Ideation process, in which the Digirati team, including UX/design experts, software developers, architects and team leads were tasked with coming up with a list of potential prototypes.

Team members worked independently to come up with proposed prototypes for exploration, each of which were written up and pitched to the entire team for collaborative discussion and eventual voting and selection.

This initial winnowing phase led to four proposals being taken forward for discussion with the Library of Congress and for the input and approval of key stakeholders.

1. Cluster and aggregate search
2. Image search
3. Fat item pages
4. Tagging and sharing

Three of these prototypes were chosen in collaboration with the Library of Congress to take forward for further development.

1. [Cluster and aggregate search](#): as an interactive prototype
2. [Fat item pages](#): as an interactive prototype
3. [Tagging and sharing](#): to be prototyped primarily as static wireframes, rather than as an interactive prototype

Pitches for the successful prototype suggestions can be found in [Appendix 1: Prototype “pitches”](#).

Once prototyping began, LoC staff were consulted regularly through the prototyping process, and each iteration of the prototypes were accompanied by show-and-tell sessions and feedback from LoC staff was incorporated into the succeeding iterations.

Technical and practical constraints

We knew before we started that there were certain practical and/or technological constraints within which we would have to work.

Infrastructure and time-to-prototype

We had approx 125 days of staff time¹ to develop the three prototypes, which had to include:

- Ideation, design and UX
- Infrastructure and DevOps
- Data retrieval and transformation
- Data enrichment
- API development
- Front end/UI development

We also had to allow for time to iterate over the prototypes in response to internal feedback and testing and feedback from key stakeholders at the Library of Congress.

With a mixed team comprising dedicated UX, front- and back-end developers, and a dedicated product manager/architect this meant that the prototypes had to be built efficiently and using as light an infrastructure as possible.

On a long-running production system: reliability and robustness, scalability, long-term data preservation, and ease of maintenance are all considerations. However, these can sometimes come with a cost in staff time—designing and provisioning the appropriate architecture—and in terms of subscription fees and compute costs for cloud-based infrastructure and software as a service.

For this project, rapid turnaround time, and low overheads in terms of staff time, were key. See [Technical approach](#) below for details of how the prototypes were deployed.

Agile prototyping

It was important that our approach had to allow us to quickly explore both data and UI/UX without committing too early to one particular approach. We knew that some of the elements in the wireframes and designs that emerged from the ideation sessions might be difficult to do well in

¹ Approx 10 weeks of development time for a blended multi-disciplinary team.

the time available. We also knew that some might turn out to be dead-ends or prove unhelpful, from a user-centered perspective, once we had initial prototypes to test against.

This meant, in practice, having loosely coupled front and backend applications with JSON APIs feeding Javascript front end applications using React and other frameworks. This made it possible to mock up static data and rapidly test frontend approaches in advance of data being available at scale. It also made it possible to work independently over multiple weeks on data retrieval and enrichment, at scale, without adversely affecting UI development.

See [Iteration 1: “Skinning the LoC JSON API”](#) and [Iteration 2: “The Fat Item Page”](#) below, for examples.

No humans in the loop

We also knew that if we wanted to explore hundreds of thousands of item records, and provide prototypes that were high-enough fidelity, in terms of their use of data, to be informative during user testing, that we could not rely on any human curation of data.

No new cataloging or data curation could happen. Instead, we had to rely on:

- Data extracted at scale from the LoC JSON API
- Data created or enriched through *unattended* machine driven enrichment processes such as entity recognition and entity linking

Prototype 1: Fat Item Pages

Problems and proposed solution

Users searching for content ultimately end up on a catalog or item page. Item pages can be quite rich sources of information but they may not always be easily comprehended by a user. Nor do they always contain all of the information that can potentially be generated about that item, or links to all of the associated sources of information or related items. This prototype was intended to address:

User Insight 3: Search and Ask a Librarian

- Users struggle to make sense of search results
- Users struggle to find what they are looking for
- User struggle to understand what the Library holds and does not
- Users fail to quickly determine whether there is an answer to their question
- Users go to Ask a Librarian when search could have answered their questions, but they were unable to find what they wanted

User Insight 4: Today Problem

- Users do not have time to spend understanding complex or difficult to comprehend resources
- Users want to quickly and effectively access content

User Insight 2: Trust

- Users want trustworthy data.
- Users want to rely on LoC to help make their information-driven decisions and arguments.

Proposal

Item pages should not just present the Marc/catalog record and a simple embedded viewer back to the user with some limited linking to related records via subject terms.

Item pages could be “fat” pages rich in sources of data to provide the user with the most information to help solve their information-driven needs, and to provide them with clear onward

journeys to likely other sources of data if this item page does not provide them with what they want.

Item pages could bring data in from other sources, such as:

- LoC collection homepage information
- LoC linked data/ authority sources, such as LCNAF
- Information from LoC research guides
- Rich crosslinking (generous interfaces)² to related records within the catalog, potentially generated via machine methods such as named entity recognition, and named entity linking
- (On the ecosystem / trust space) we might link out to other records in related US governmental sources like archive.gov (which has a JSON API, so this could be done)
- Potential 3rd party sources such as Wikidata or wikipedia
- Presenting summary information and aggregation information about related records

Even if this results in repetition across catalog pages, the beginner user, or the person beset by the “Today problem” isn’t likely to look at many item pages, and we hypothesized that they would appreciate the rich information being presented to them (with careful UX treatment).

Process: Wireframing, design and further ideation

Before we could begin active development on this prototype we wanted to ensure that the entire team had a chance to contribute their thoughts and insights in terms of the features of the prototype and the visual design and user experience that the prototype should embody.

Our process for each prototype was to, independently (“working alone together”):

1. Review the concept card / pitch deck
2. Carry out independent desk research. For example, identify examples of both good and bad attempts to solve similar problems elsewhere or in the past.
3. Review the JSON API data from the LoC to assess what data might be available to assist in prototyping
4. Sketching

Then, collectively (as a team):

² <http://www.digitalhumanities.org/dhq/vol/9/1/000205/000205.html> and <https://pro.europeana.eu/page/issue-11-generous-interfaces>

5. Vote, in order to identify which sketches or individual features of particular sketches met with the widest approval across the team, with success in addressing the identified How Might We problems as the primary criterion
6. Consolidate the highest ranked features into a set of consolidating/unified wireframes
7. Create the product backlog and user stories for prototype development

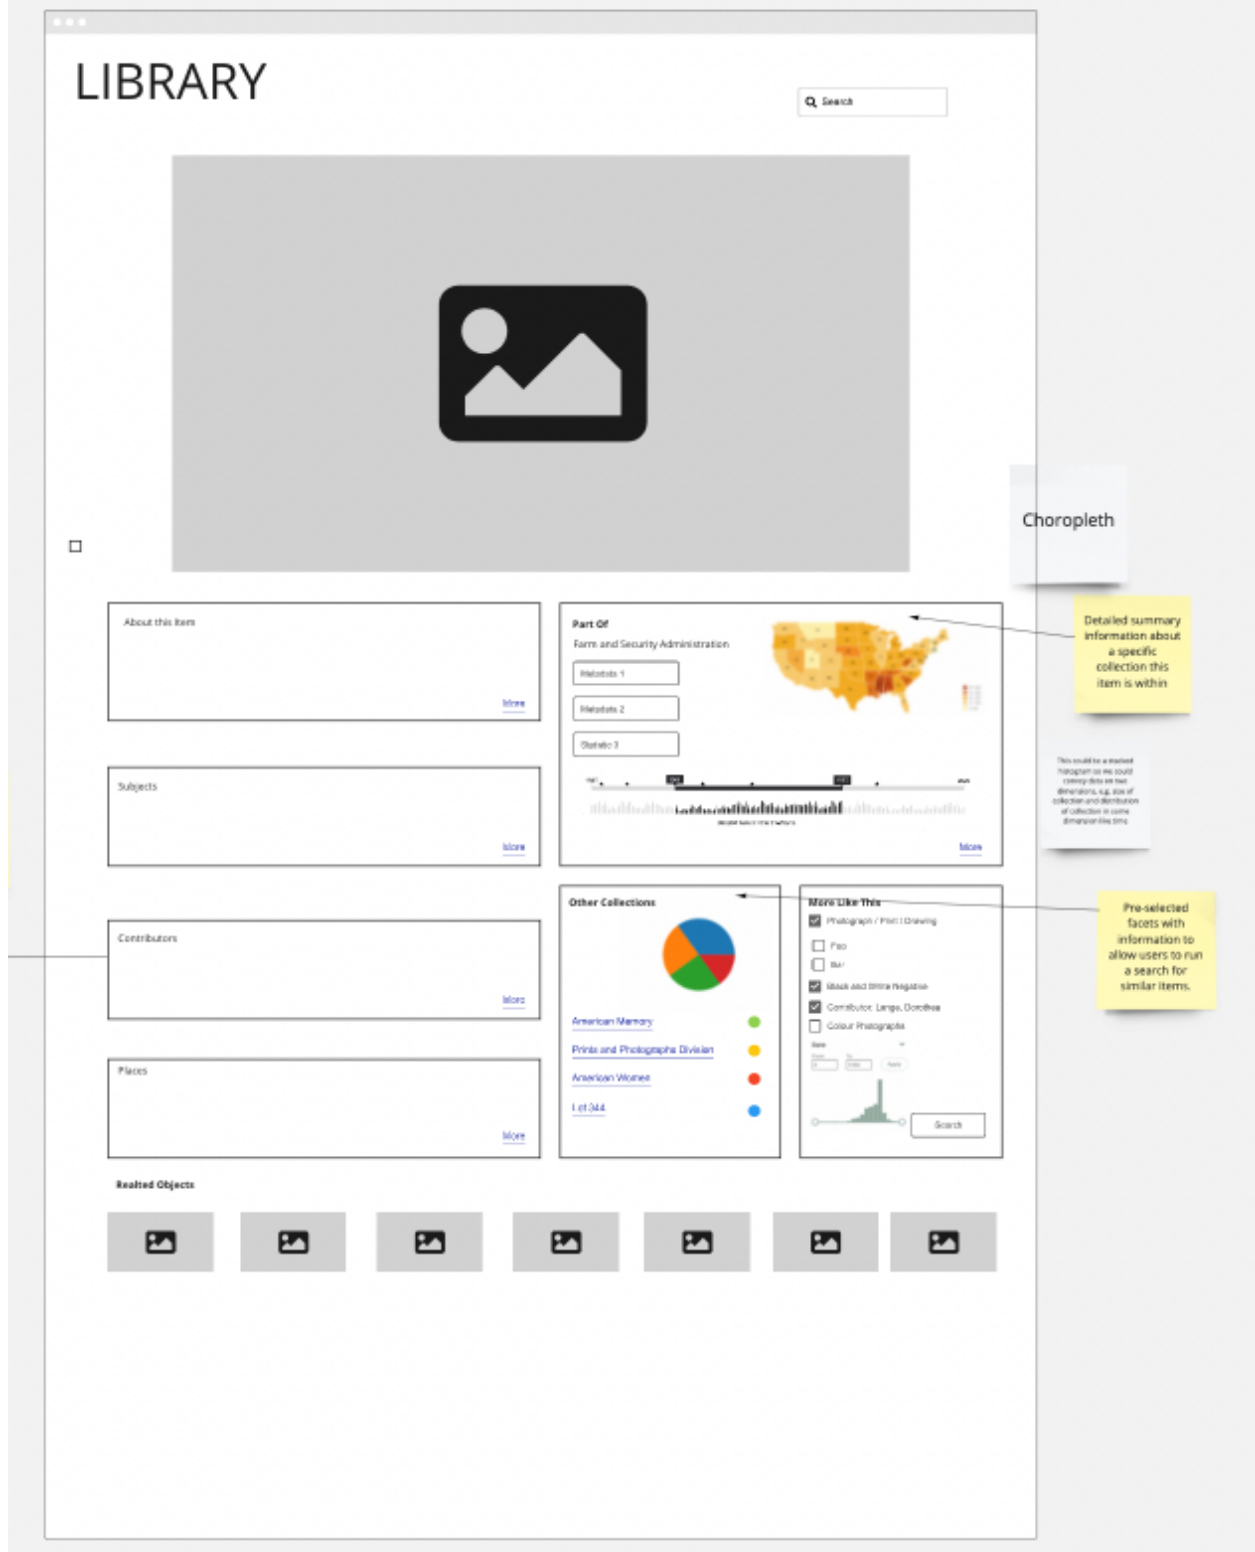
LoC staff had opportunities to feed back on this work once the prototype development began.

User Stories

The (abbreviated) user stories that both emerged from and informed the sketches included:

- As a user I can view this object online ...
- As a user I understand how this record fits into the wider data landscape ...
- As a user I want to know what I might find when I click on a link ...
- As a user I can find information that puts this record in context ...
- As a user I want information that helps me quickly understand what I am looking at ...
- As a user I can see information about the people, places and organizations related to this record ...

Multi Column Layout



Example of an annotated sketch for the Fat Item Page

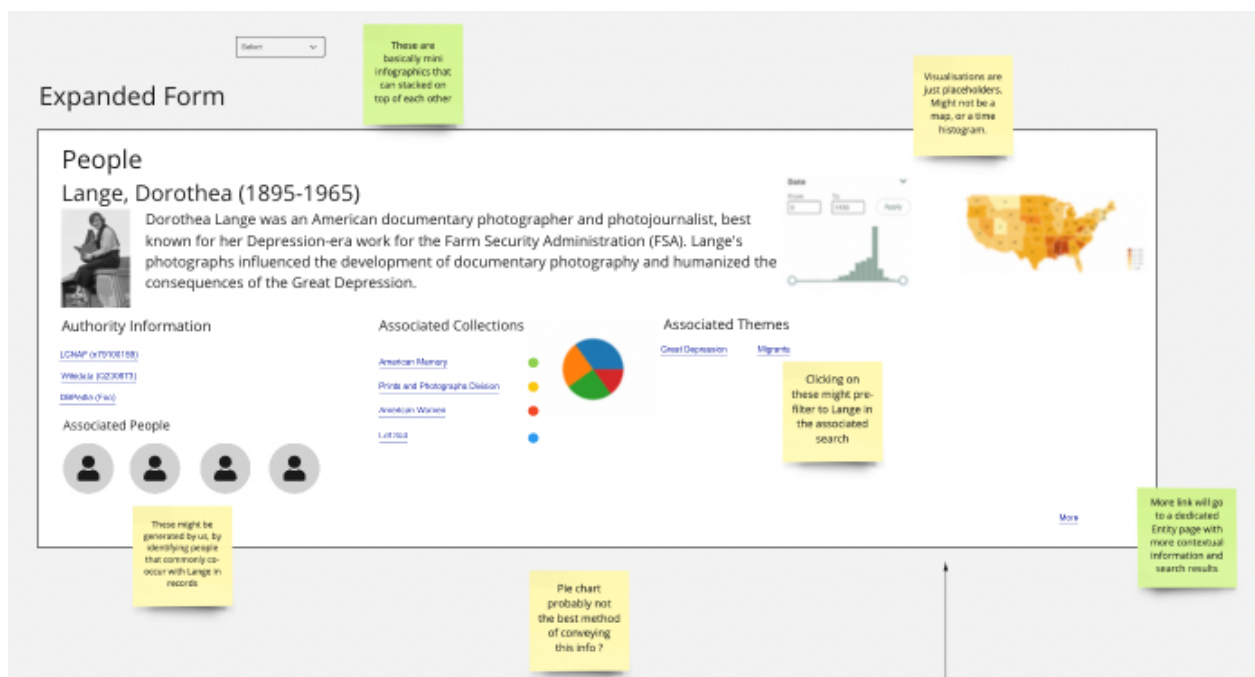
Design/UX features

There were a number of features which emerged from the sketching sessions which we wanted to take into the prototype:

Popups/previews

The information available related to the object on the page might, potentially, be quite rich, and include information about *entities* (people, places, other items that might relate to some controlled vocabulary item or other). Ideally, it should be possible to convey some of this information in the page, either via an informative pop-up or a preview embedded in some way within the page, so that the user:

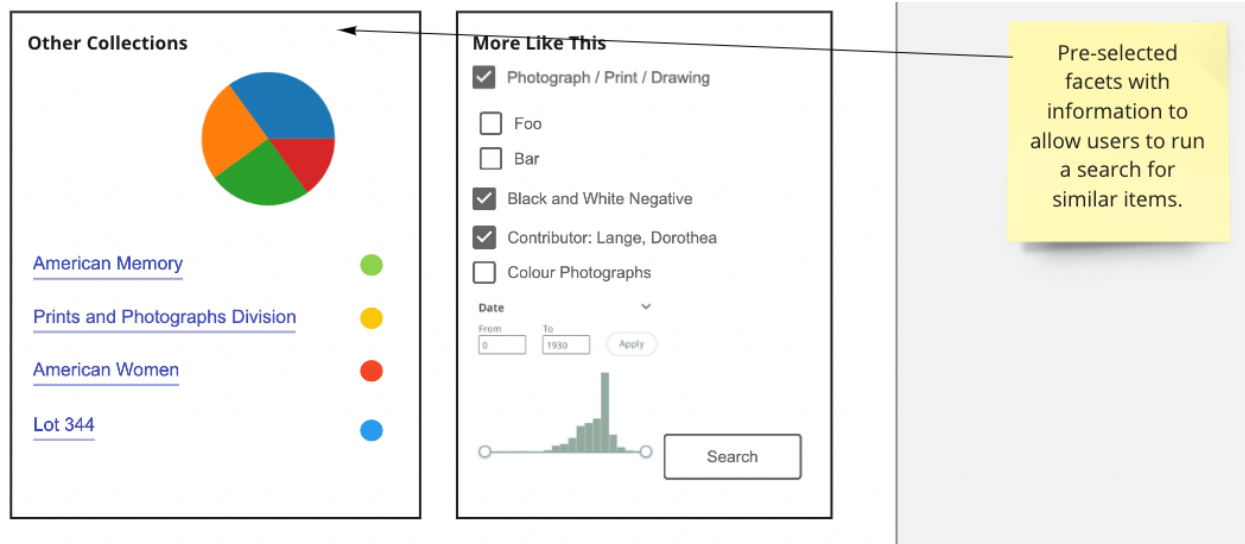
- has additional contextual information to understand this record
- Is not forced to click away from the page in order to find additional information



Annotated sketch from pre-development ideation workshops

Interactive facets for similar objects

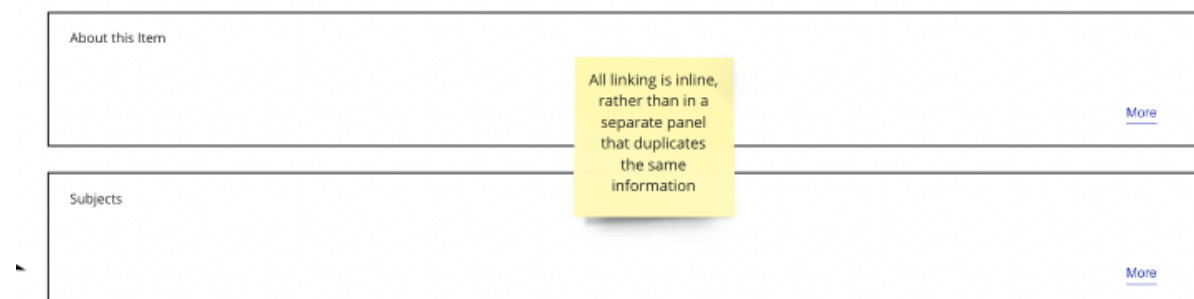
Similarly, users may wish to understand—perhaps view a preview of faceted search over related objects—what related or similar objects are available without clicking away from the item currently being viewed.



Annotated sketch from pre-development ideation workshops

Generous interfaces / embedded links

The current LoC interface provides onward links via a separate panel on the page. The bulk of the content is passive text without any interlinking with metadata terms in descriptive fields or links through to searches for related or similar objects. We felt that embedding links within the text of the catalog record would make for a more generous³ and approachable user interface.

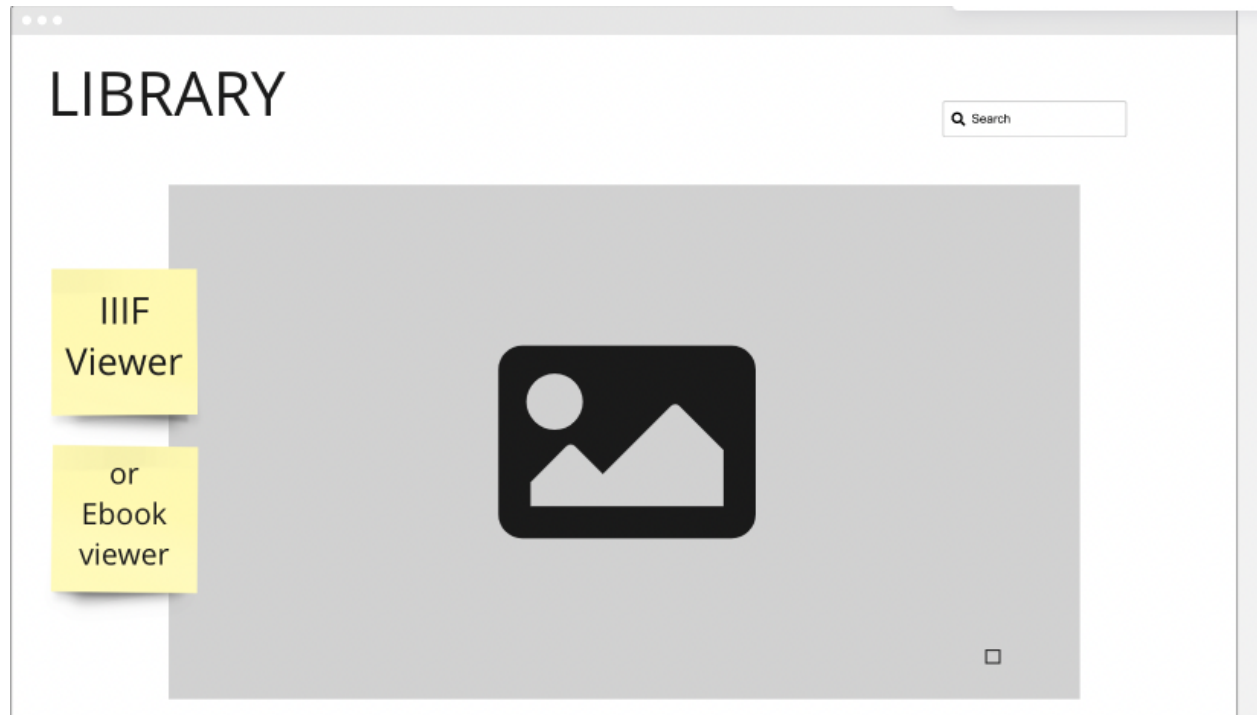


Annotated sketch from pre-development ideation workshops

³ <http://www.digitalhumanities.org/dhq/vol/9/1/000205/000205.html> and <https://pro.europeana.eu/page/issue-11-generous-interfaces>

Embedded IIIF based viewer

Another key element of our vision for the Fat Item page was that the viewer should be embedded within the page. The user should be able to explore digitized content, where that content exists, within the context in which that content is discovered and described.

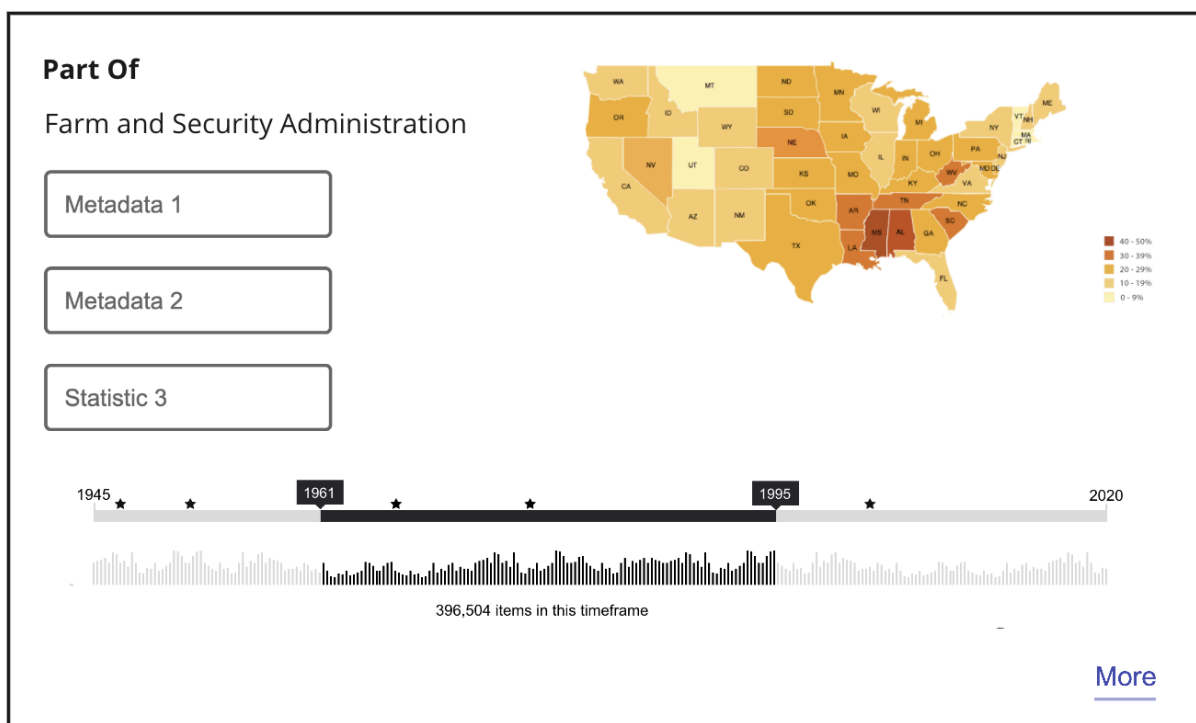


Annotated sketch from pre-development ideation workshops

Use of maps and/or timelines

Many of the objects in the LoC collection can be located in space and time. Either via their date or place of creation, or via places or dates that are part of the subject matter or content of that collection item.

Presenting spatial and temporal information in a visual manner was another one of the elements that fed into our understanding of how we might enrich ("fatten") the item page.



Annotated sketch from pre-development ideation workshops

Development

While the “Fat Item Page” prototype went through many small incremental iterations, there were two key separate prototype phases which adopted different approaches to the data.

Development on the two main interactive prototypes—Fat Item pages, and Aggregating and Clustering—were not carried out consecutively, or in isolation from each other. Instead, the two prototypes were interleaved, with effectively four key phases:

1. Item Page 1: [Skinning the LoC JSON API](#)
2. Data retrieval, enrichment and APIs (see [Enrichment](#) below)
3. Item Page 2: [The Fat Item Page](#)
4. [Aggregation and Clustering](#)


Iteration 1: "Skinning the LoC JSON API"

The first version of the Fat Item page that we built relied on the existing LoC JSON API to provide the skeleton of the data to the frontend, and also leveraged the existing LoC JSON API to drive search across objects and onward linking to related objects.

In this version we used the existing LoC JSON API for item records to drive a frontend prototype in which we could test some of the elements from our sketches and wireframes (see [Process: Wireframing, design and further ideation](#) above).

PHOTO, PRINT, DRAWING

Lorraine Motel, Memphis, Tennessee



- About this item
 - Part of
 - Facets
- Rights & Access
- Cite This Item
- JSON Response
- More like this
- You might also like

About this item


Title
Lorraine Motel, Memphis, Tennessee

Other title

Contributor Names
[Highsmith, Carol M.](#), 1946-, photographer

Created / Published
2006 April 17.

Subject Headings
- United States--[Tennessee](#)--[Memphis](#)
- Lorraine Motel
- Motels
- [Martin Luther King, Jr](#)



Item page based on the existing LoC JSON API

See: <https://loc.ch.digtest.co.uk/site/item/2010630679> for the above page.

We were able to test:

- Embedding a viewer
- Using maps to locate objects in space
- Generous linking from within the catalog item text
- Interactive facets within the Item page for onward journeys/search

Find similar items

Format

☐ Photo, Print, Drawing

Contributors

☐ Highsmith, Carol M.

Dates

☐ 2006

Locations

☐ Memphis

☐ Tennessee

☐ United States

Languages

☐ English

Subjects

☐ America

☐ Assassination

☐ Color

☐ Digital Photographs

☐ Lorraine Motel

☐ Martin Luther King, Jr

☐ Martin Luther King, Jr.

☐ Memphis

☐ Motels

☐ Tennessee

☐ United States

[Go back](#)

Facets on the Item Page

This approach let us rapidly test some of our UI elements before we had completed the backend work harvesting and enriching LoC catalog records.

In addition, we were able to test *enriching* these records via entity linking *on-the-fly*, such that the existing LoC JSON API data was decorated with new, machine-driven additional data built from that catalog record data.

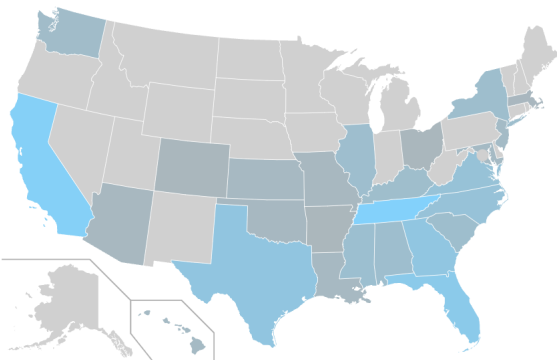
This prototype even linked to a search page, decorated with maps and an improved facet UI, driven by the existing LoC JSON API

LIBRARY

UNIVERSITY OF CONGRESS

EXPERIMENTAL PROTOTYPE

View on map




Results: 1-100 of 7689

Refined by:

Available Online (x)

Location: tennessee (x)

List




MAP

Battle of Lookout Mountain, Tennessee November 24th, 1863.

Oriented with the Tennessee River to the top and the East Tennessee Railroad to the east, image gives details of the difficult terrain surrounding the Confederate position on Lookout Mountain. Dotted lines show the directions of the Union attack and the Confederate retreat.

Date: 1861



MAP

Sanborn Fire Insurance Map from Rockwood, Roane County, Tennessee.

May 1903. 5.

Date: 1903-05

Available Online

True

False

7689

15998

Original Format

Photo, Print, Drawing

Personal Narrative

Map

Book

Manuscript/Mixed Material

Newspaper

Notated Music

3906

1368

992

376

323

248

199

See: <https://loc.ch.digtest.co.uk/site/search/?fa=location:tennessee&fo=json> for the above search.

This prototype was an effective way to test some of the UI components and the general approach to the Item page, without committing to extensive backend development. Many of the ideas that emerged from this process were carried forward into the next stage of development.

30

Iteration 2: “The Fat Item Page”

Not all of the features we wanted to test for the Fat Item Page could be delivered solely on top of the existing LoC JSON API.


The second iteration of the page made much more heavy use of the [enrichment](#) of data through entity linking, and also leveraged the search and other APIs developed for the [Aggregation and Clustering](#) prototype to provide the generous interfaces and onward linking. Our use of entities was restricted to concrete people, places and organizations, and did not include themes, subjects or other abstract concepts.

This page also made use of a Node.js based proxy/transformer to transform non-IIIF resources provided from the LoC APIs into IIIF manifests for embedded viewing.⁴

PHOTO, PRINT, DRAWING

Lorraine Motel, Memphis, Tennessee

original digital file



1 of 1 • Image 0

- About this item
 - Part of
 - Facets
- People
- Locations
- Rights & Access
- JSON Response

About this item

Title
Lorraine Motel, Memphis, Tennessee

Contributor Names
Highsmith, Carol M., 1946-, photographer

Subjects

- United States--Tennessee--Memphis
- Lorraine Motel
- Motels
- Martin Luther King, Jr
- Assassination
- America

Notes

- Lorraine Motel, where Dr. Martin Luther King was assassinated during the Civil Rights Movement on April 4, 1968.

Part of What is this?
Carol M Highsmith

Find similar items What is this?

Language
English

Original format
Photo, Print, Drawing

Subjects

- United States--Tennessee--Memphis
- Lorraine Motel
- Motels
- Martin Luther King, Jr

⁴ See: <https://loc.ch.digtest.co.uk/site/items/24013891> for an example with a full multi-page manifest.

Popups and previews

This prototype implemented popups and previews to show users information about linked data.

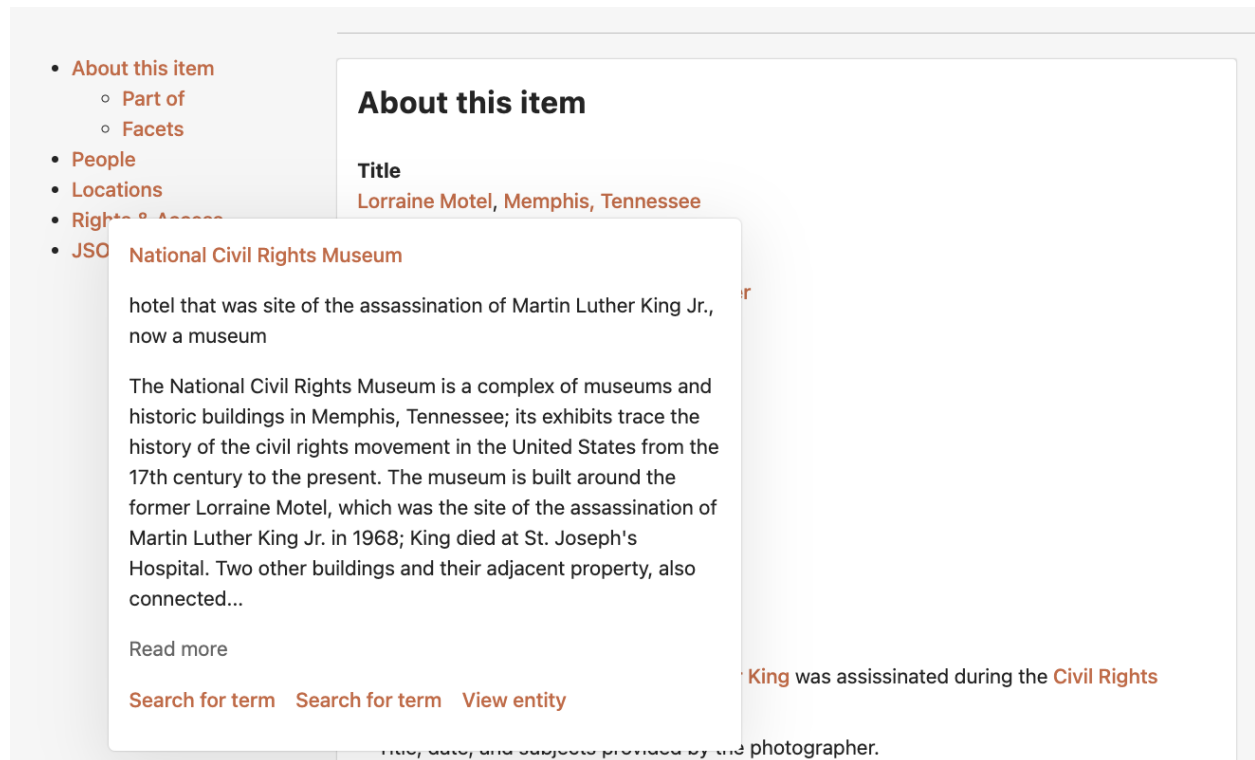


Image showing the popup when a user clicks on Lorraine Motel

These pop ups provide onward links to search, but also to dedicated *Entity Pages*⁵ which bring together multiple sources of information about a particular person, place, or organization

Where this information exists, some of that information can be brought back into the Item page. See below for an example of Wikipedia information in an Item page popup.

⁵ See <https://loc.ch.digtest.co.uk/site/entities/wikidata:Q60> for an example entity page.

About this item

Title

Lorraine Motel, Memphis, Tennessee



Memphis

county seat of Shelby County, Tennessee, United States

Memphis is a city along the Mississippi River in southwestern Shelby County, Tennessee, United States. Its 2020 population was 633,104, making it Tennessee's second-most populous city behind Nashville, the nation's 28th-largest, and the largest city proper situated along the Mississippi River. Greater Memphis is the 42nd-largest metropolitan area in the United States, with a population of 1,348,260 in 2017. The city is the anchor of West Tennessee...


[Read more](#)

[View entity](#)

Wikipedia preview information for Memphis, Tennessee

Entity information

Embedded within the page is rich information about people, places and organizations.


 **People** (External content)

What is this?

Carol M. Highsmith

Carol McKinney Highsmith (born Carol Louise McKinney on May 18, 1946) is an American photographer, author, and publisher who has photographed in all the states of the United States, as well as the District of Columbia, and Puerto Rico. She photographs the entire American vista (including landscapes, architecture, urban and rural life, and people in their work environments) in all fifty U.S. states as a record of the early 21st century. Highsmith ...


[Read more](#)



Martin Luther King Jr.

Martin Luther King Jr. (born Michael King Jr.; January 15, 1929 – April 4, 1968) was an American Baptist minister and activist who became the most visible spokesman and leader in the American civil rights movement from 1955 until his assassination in 1968. King advanced civil rights through nonviolence and civil disobedience, inspired by his Christian beliefs and the nonviolent activism of Mahatma Gandhi. He was the son of early civil rights acti...

[Read more](#)



People information within an [item page](#)

The overall approach—addressing the insights from user research, see [Problems and proposed solution](#) above— is to leverage as much of the catalog data and data from machine enrichment as possible, to provide as rich and “fat” as possible an experience for the user without the user leaving the page.

Technical approach

The broad technical approach for the [Fat Item Page](#) and [Aggregation and Clustering](#) followed the [Agile prototyping](#) model above.

The page consists of a lightweight [React](#) based UI provisioned by a dedicated “Fat Item API” which provides all of the data required by the frontend in a single JSON payload.

See: <https://loc.ch.digtest.co.uk/api/loc/data/item/2010630679> for an example Item API return, which includes:

- The original LoC Resource JSON data from the LoC API
- Entity data for people, places and organizations
- Highlight information to allow for the creation of links and the provisioning of popups within the text.
- Search data to provide the UI with the correct links to call search when a user follows a link
- Facet data for the sidebar
- A link to a stored/cached copy of a IIIF manifest generated from the LoC JSON API data

Given the constraints in terms of time, and especially in terms of time spent on infrastructure and DevOps (rather than on development and user experience), we adopted a relatively lightweight approach to architecture and infrastructure:

1. Amazon AWS [EC2](#) (Elastic Compute) and [ECS](#) (Elastic Container Service) were used for hosting the core services.
2. Core services were containerised so that we could easily run a local instance of the stack for development and testing
3. Amazon [RDS](#) was used to host the [PostgreSQL](#) database we used to store the data and provision the APIs
4. Continuous integration and continuous deployment was provisioned via [Github actions](#). Updates would automatically deploy to AWS once a commit was merged into a protected branch and all tests had passed.

We ruled out using *at run-time* any cloud-based services that we could not easily run locally, so, for example, queues for long-running asynchronous processes (such as data retrieval and enrichment) used queue brokers and tooling that could be run easily in a local containerised development environment, rather than leveraging enterprise level workflow/stream processing tools or cloud-based services.

User insights

'There are rich opportunities for me as an instructor to have students start with this unambiguous image. I can use this page as a national teaching instrument as it is.'

Digital learning design specialist, former HS Teacher

Information-dense pages are welcome

The main idea behind this prototype is about enriching the catalog items. This was perhaps the concept that got the users most engaged. The ability to have a catalog item that serves as a platform to discover more connections had users excited. They felt this type of page can give enough information at different levels and wouldn't overwhelm beginners.

People, locations: 'incredible useful' for educators and students

During these sessions, we were able to confirm how enriching catalog items with people and locations can get more people interested in the materials at LoC. Those we spoke to that work in education assured us that these pages could be useful resources for educators and students.

One user said:

Other human stories resonate with people. Archival object data not that much.

Provenance matters to people

As mentioned by members of the LoC team earlier during the conception of these prototypes, people do want to know about the provenance of data. They understand that to benefit from an enriched experience there would be data from external sources. Users wanted to clearly understand the provenance of all data in a page:

Most people are fine with Wikipedia... This has a different source and I can take it in strides. I can either fully qualify it or have a different value for that. But to me, it doesn't detract.

Citation

We have been told students (and teachers) love citing but many don't know how to do this properly. The prototypes didn't include a citation widget such as the one in the LoC website. But based on what we have seen in these sessions people are missing the Rights and access info component entirely, a similar UI component to Citation, LoC might consider exploring further how both of these components could be made more visible and easier to spot.

Usability notes

Main image

The main image's zoom controls got in the way for some users that wanted to scroll down with their laptops.

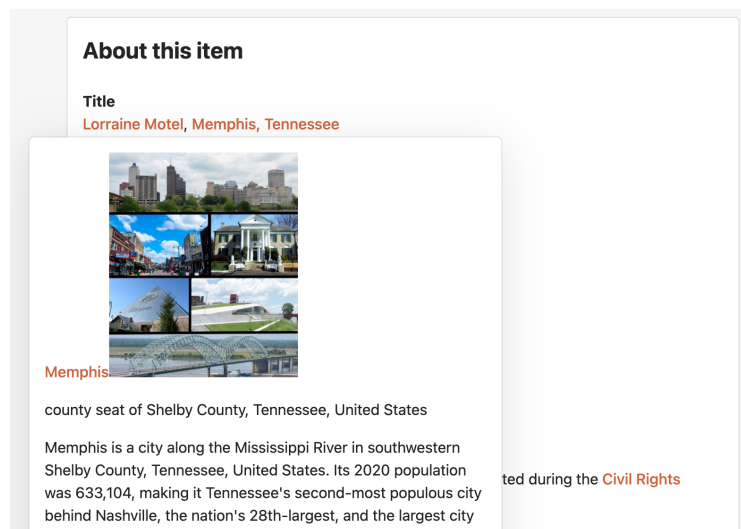
The image menu was discussed briefly for the few that spotted it and interacted with it. They weren't sure about the options available there. There were a couple of bugs too: the image menu showing twice, on each side of the image at some point. Also, buttons for previous and next images are shown but this was just a single image.

People and location

This section, in particular People, got users more engaged than the archival information in this page. Users commented on the fact they knew the picture but didn't know the photographer - Carol M. Highsmith. This got them more interested in her work as a couple of users went their own way to explore more about her. One user summed up the general view:

Other human stories resonate with people. Archival object data not that much.

Popup preview



Addressing the insight mentioned in the search section earlier: *clicking on an item is a commitment*, this was well received all across. It might boil down to the fact that people can remain on the page they have an interest in and explore just enough on other items without losing focus on their target.

The preview should be triggered on hover (rollover effect) rather than on click though. Interestingly, a few people also added that they tend to open items in another tab fearing to lose what they were looking for initially.

Find similar items

This was the hardest component for users to understand and use. Initially they saw it as an extension of related metadata information without much to add. All missed the *Combine* button at the bottom. But once found, the functionality became clear to all. Placing the button also at the top was suggested by just about everyone. Additionally, we believe this would need to be tested more thoroughly, potentially exploring a simplification that would display the checkboxes upfront to be more in line with people's mental model for this type of interaction.

Navigation and layout

Users managed very well the density of the page and they recognised both the content broken into clear chunks/sections and the section menu (located on the right hand side) as the key elements that made their experience positive.

They didn't mind scrolling down as it 'felt natural' to them and the content wasn't 'all jammed together'. They commented they could see this working well on mobile too.

Rights and access

A few users had a sense of urgency with regards to image use rights. It wasn't clear to them where to find this information in the page. They even suggested having a search filter for this. They didn't notice the Rights and Access component which was quite down the page.

An additional single clear label about copyright very close to the image might solve this issue.

Prototype 2: Aggregation and Clustering

Problems and proposed solution

This particular set of problems impacts on all of the core user types. Carrying out searches and understanding the results of those searches is a fundamental requirement for almost every user in their journey to discover and use the Library's collections. This requirement is particularly difficult to deal with when users are either inexperienced at using and understanding library catalogs, or when users are under time pressure.

User Insight 3: Search and Ask a Librarian

- Users struggle to make sense of search results
- Users struggle to find what they are looking for
- User struggle to understand what the Library holds and does not
- Users fail to quickly determine whether there is an answer to their question
- Users go to Ask a Librarian when search could have answered their questions, but they were unable to find what they wanted

User Insight 4: Today Problem

- Users do not have time to spend understanding complex or difficult to comprehend resources
- Collections are not optimally organized for time-pressured users
- Users want to quickly and effectively access content

Proposal

The data in the catalog can be understood as a rich multidimensional nested graph of information, so we could explore multidimensional ways to represent this rich graph-like data, rather than as a unidimensional flat list.

So that a user can browse/navigate to the group of results they want and then from there to more specific clusters and individual records, instead of presenting the search results as a single ordered list with facets, present the information:

1. Clustered hierarchically into groups
2. With useful summary information about each group in aggregate
3. Employing techniques from data visualization to make the results easier to comprehend at a glance

Process: Wireframing, design and further ideation

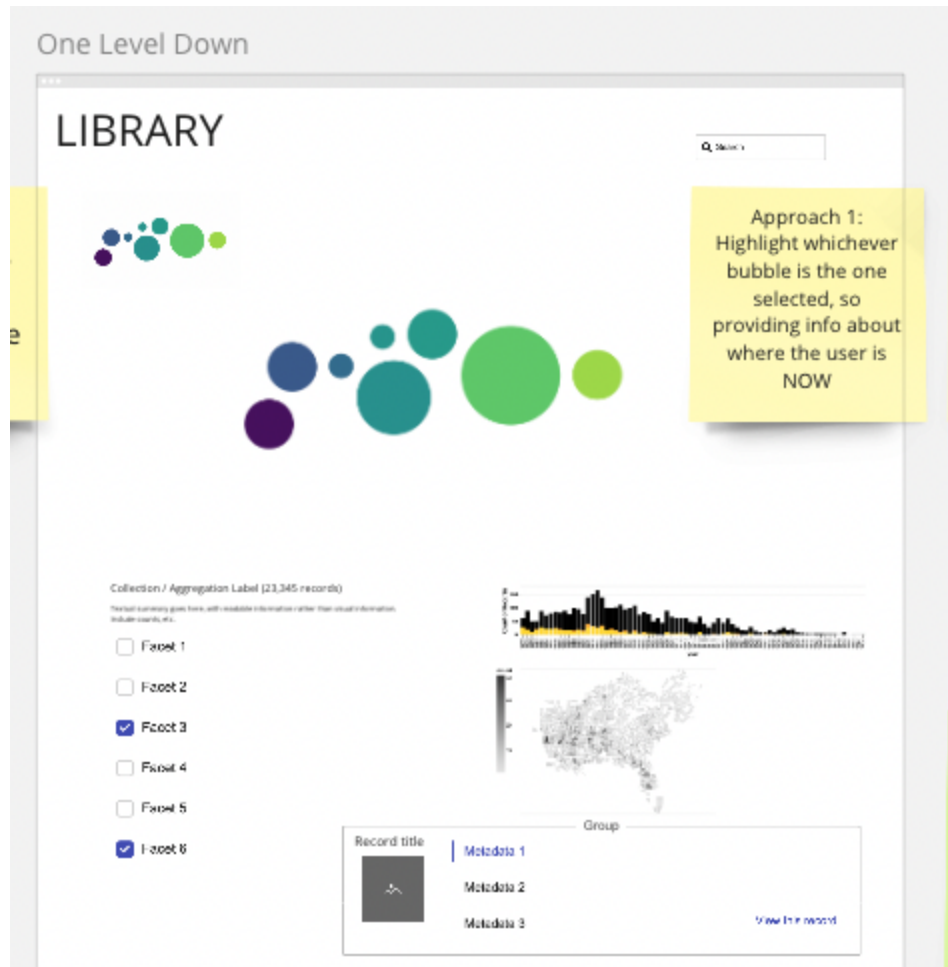
As with the Fat Item Page prototype, before we could begin active development on this prototype we wanted to ensure that the entire team had a chance to contribute their thoughts and insights in terms of the features of the prototype and the visual design and user experience that the prototype should embody.

Our process for each prototype was to, independently (“working alone together”):

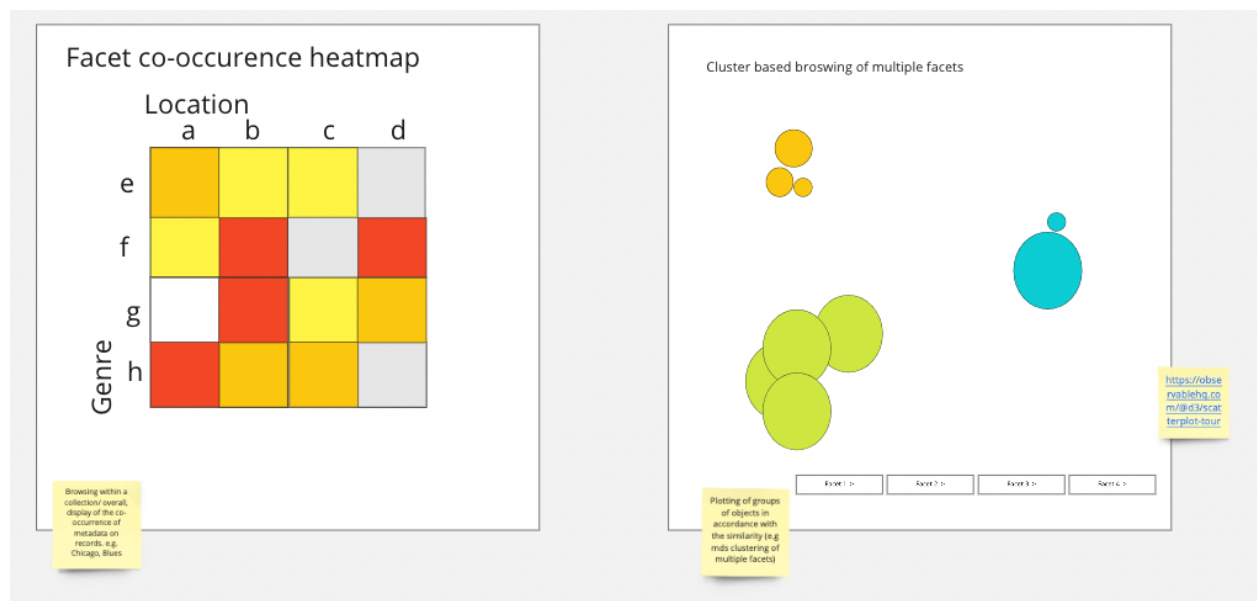
1. Review the concept card / pitch deck
2. Carry out independent desk research. For example, identify examples of both good and bad attempts to solve similar problems elsewhere or in the past.
3. Review the JSON API data from the LoC to assess what data might be available to assist in prototyping
4. Sketching

Then, collectively (as a team):

5. Vote, in order to identify which sketches or individual features of particular sketches met with the widest approval across the team, with success in addressing the identified How Might We problems as the primary criterion
6. Consolidate the highest ranked features into a set of consolidating/unified wireframes
7. Create the product backlog and user stories for prototype development



Sketch of one initial approach to clustering and aggregation

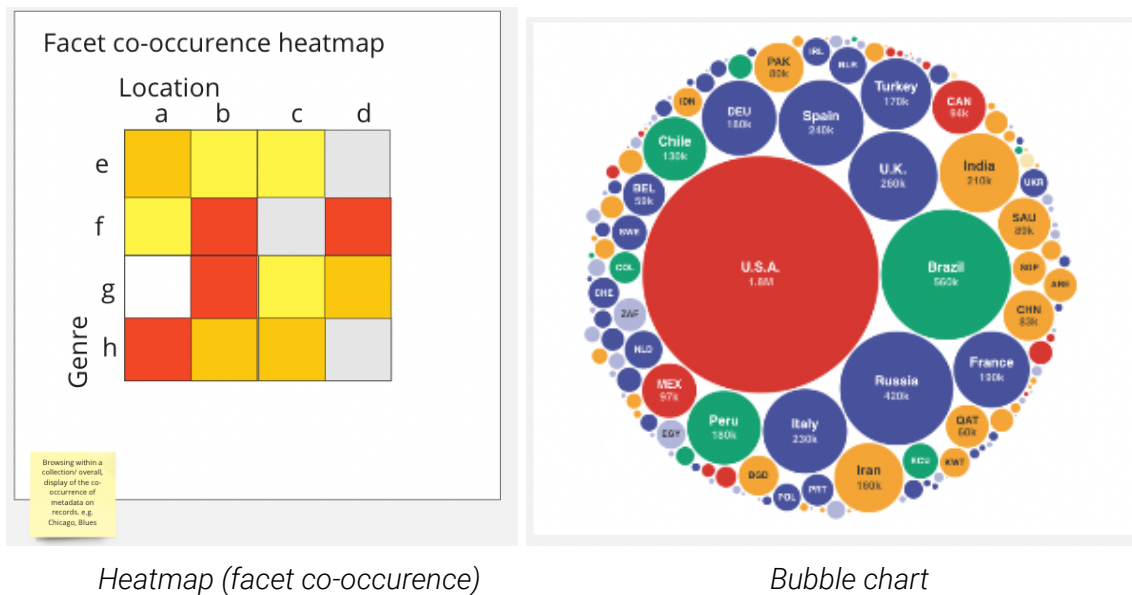


Two different approaches to visualizing multiple facets at one time

Design/UX features

There were a number of features which emerged from the sketching sessions which we wanted to take into the prototype.

Visualization of metadata or entity driven facets

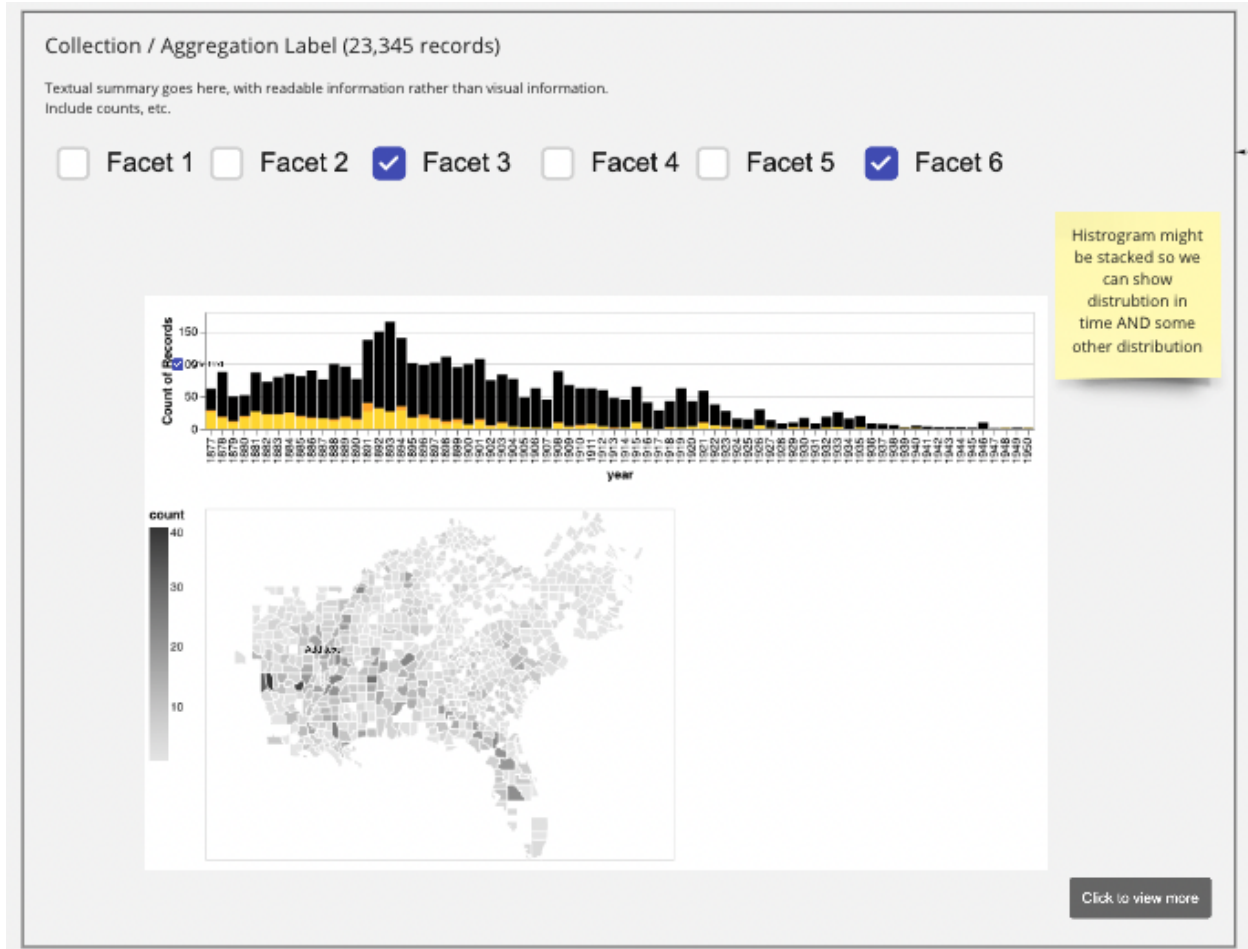


Visual navigation

Another idea that emerged during initial sketching was the idea that visualizations should not be passive, but instead, be a way of further filtering results or navigating through a collection.

Summary information

Provide the user with information about how some grouping is distributed in time or space, or about other topics and/or entities that exist alongside this one.



Distribution in time and space (sketch from ideation workshops)

Development process

While the first prototype went through many small incremental iterations, the *Aggregation and Clustering* prototype was built initially via many small standalone front-end components, which could be used to test individual pieces of UI/UX in isolation.

The final prototype was then assembled once the Enrichment and API work (see [Enrichment and APIs](#) below) was finished and we had a good understanding of the shape of data we would require for the front end.

One approach that was very effective in the initial development of front-end visual components for the *Aggregation and Clustering* prototype was to create a mixture of static and dynamic data (as JSON) and then feed these into lightweight components hosted in [CodeSandbox](#).

We also made use of lightweight React based wrappers around more complex visualization libraries, such as [D3.js](#) and [Semiotic](#). This made it possible for us to democratize the development of frontend components across the team and test a wide range of options.

As we narrowed down on promising choices we were able to hook the CodeSandbox hosted prototypes to real public APIs on our QA instance and test with live data.

Code Sandbox Examples

- [Timeline histogram](#)
- [Co-occurrence matrix with API hookup](#)
- [Network diagram 1 with scaled bubbles](#)
- [Network diagram 2 with fixed bubbles](#)
- [Network diagram \(different force layout\)](#)
- [Scatter plot 1 with items](#)
- [Scatter plot 2 with facets only](#)
- [Large co-occurrence matrix](#)

Approaches we ruled out

Dimensionality reduction and K-Means clustering


This was promising, in terms of generating synthetic clusters out of the data, but was too slow to do in real time. However, this might be something that could be explored further in a future phase by pre-computing data and then visualizing that data. With relatively unchanging data, such as a collection of catalog records, it ought to be possible to generate useful clustered data via longer running asynchronous processing.

Facet co-occurrence matrix

This was also very promising, and was prototyped extensively in CodeSanbox.

Approaches we kept in

- Facets clustering using a force-directed layout to group co-occurring facets together
- Chord diagram showing the facets in combination
- Network diagram showing facets AND items together



See: <https://loc.ch.digtest.co.uk/site/new-search?q=luther> to view this search and the associated visualizations on the prototype site.

In addition, on the pages we created for *Entities* (see below) we also kept in the ideal of visualizing the distribution of collections in space. For example on the entity page for Martin Luther King (<https://loc.ch.digtest.co.uk/site/entities/wikidata:Q8027>) we can see the spatial distribution of that subset of collection items tagged (by enrichment) with Martin Luther King.

The map element here was produced using SVGs, rather than using shapefiles on a map (for example using a Leaflet based viewer) for performance reasons, and worked well for dynamically generating lightweight performant maps.


Martin Luther King Jr.



American clergyman, activist, and leader in the American Civil Rights Movement

Martin Luther King Jr. (born Michael King Jr.; January 15, 1929 – April 4, 1968) was an American Baptist minister and activist who became the most visible spokesman and leader in the American civil rights movement from 1955 until his assassination in 1968. King advanced civil rights through nonviolence and civil disobedience, inspired by his Christian beliefs and the nonviolent activism of Mahatma Gandhi. He was the son of early civil rights activist and minister Martin Luther King Sr. King participated in and led marches for blacks' right to vote, desegregation, labor rights, and other basic civil rights.

Geographic distribution of catalogue items



4 RESULTS

Search

Featured in

Carol M Highsmith

This topic appears 50 times in this collection

Contributor Names

Highsmith, Carol M., 1946-, Photographer

4

Created Published Date

2010 February 28.

2

2006 April 17.

1

2010 April 21.

1

Entities

Location: United States Of America

4

Collection items for "Martin Luther King" showing map

The build of the final *Entity* and *Search* pages (where we implemented the *Aggregation and Clustering* prototype elements) were dependent upon the provision of APIs for entities and search, and on the enrichment of the catalog records using machine methods. See the next section for how this was done.

User insights

'You need to make decisions on the spot. Like a music DJ: 'am I gonna play this record for the crowd or that record for the crowd?'

Investigative journalist, musician, author

Search behaviors

All users interviewed spend considerable amounts of time using the built-in search functionality within the sites that hold the resources they are interested in. Although some of the users mentioned that they go back and forth between Google and the site to help pinpoint resources that the site are not returning for them in the search as expected - they believe they should be there somewhere, as one user put it:

The problem with LoC material is that it's hard to find. You need to dig deep. Needs constant cross-reference with Google to find what you're looking for. Takes time to find.

Personalization

One aspect the exploration of this concept brought up is the expectation of users to see content that matters to them. Both people in the arts and in education mentioned how important it is for the content, its organization and labeling to work in tandem to deliver truly personalized results. They link this expectation with the ability to find the content they are after quicker as well, which is an insight first mentioned in our discovery report.

If somehow the system knew I came as an educator this would be more accessible to me. It looks like the work hasn't been done to put a filter taxonomy that would help me and filter based on how I am going to use these materials.

As in the quote above, the use of a **taxonomy** within LoC was mentioned as a way to help users behind the scenes. On the value they found on some of the filter categories, another user added:

Most people won't know the publishing date.

Curation

Related to personalization is curation. Tech savvy users wondered whether IA could play a part here. From identifying what type of user a person is early in their interactions with LoC all the way to serving them with results that 'matter to them' as one user put it.

Clicking on an item is a commitment

Users want sufficient enough information to provide them context on what they could be about to see before committing to click on an item. See sections 'Search results lists' and 'Popup preview' for more information on this.

If I could get more information prior to committing to another page then I can feel inclined to or not to - which is also successful for me.

Labels for all

This point is beyond our study here but everyone, including people that work as research librarians, felt at odds with quite a few categories (labels) in the filters. Everyone tried to help and came up with their own substitute labels. LoC might want to consider **card sorting** exercises with various audiences to explore this in more detail. It may well feed into the personalisation insight mentioned earlier.

This is a challenging filter for teachers and students. The filter types are classic archivist domain terms. We work with similar agencies, we tell them you made a wonderful tool for you.

We provide more examples of this topic in the usability section when discussing filters.

Search vs. discovery

The addition of the interactive data visualizations in the search prototype introduced the possibility to explore relationships in the data further. Users in education-related roles highlighted that students tend to be more open to explore but teachers are under tighter time pressures and have more focused tasks at hand so accessing what they are looking for quicker is more important to them. Similarly, those using LoC collections for creative and journalist endeavors were very clear about getting the content they need first and exploring second.

You need to make decisions on the spot. Like a music DJ 'am I gonna play this record for the crowd or that record for the crowd?

Usability notes

Interactive visualizations



This is an area we wanted to explore with interest. Because of that, we made the design decision to expose the visualizations high up in the user interface so people didn't miss them. People did question the placement above the search results. One user said:

Why are these before the results? I imagine there might be a good reason and I would explore these. I'd like more information about these before I abandon my search.

Users recognised that these could be very useful for early stages of research and discovery and are perhaps most appreciated by visual thinkers - the research librarians we spoke to called them 'nice artwork' and 'colorful whatevers'.

We believe the lack of some default labels - as opposed to labels that only show when interacting with the visualizations - made these a bit intimidating for some.

Interactive visualizations can easily be considered a main feature on their own and we recommend exposing more users to these for feedback. We would suggest working on the user experience of these in more detail, bringing elements of information architecture, interaction and visual design together in a format to ensure users who might not be entirely familiar with this type of data representations feel curious rather than daunted by them.

Search results list

Users explored this area with interest. The fact that some items had images made some users wonder whether there could be more descriptive texts as well (on some items) so they can make a more informed and quicker decision on committing to the next page.

Filters

Everyone we spoke to is very familiar with the concept of filters. They expected them to be on the left hand side of the page as opposed to the right as per the prototype. Though the biggest issue people encountered was with the categories listed and their perceived usefulness. There was a fair level of agreement amongst users that most of these are of no interest to them - they were quite strong about this. They appreciate that these would be of more interest to librarians and archivists but they didn't find value in them.

Entities was a label that confused many users in these sessions. They weren't quite sure what that meant. Even a research librarian was trying to suggest alternative labels for it.

However, one category, which happened to be further down in the filters, got people excited: *Subjects*. They immediately saw how that could be beneficial to their own searches. They also expected categories to be more in line with the subject a user was researching. For example, while searching on Martin Luther King, they wanted to see filters organized by topics, era, events in his life, etc.

Users also felt the bars used to visualize the volume of results was a nice addition but they didn't understand its purpose initially. They also wondered how it would look with more data as the example explored had a low number of items visually represented.

Created Published Date



Enrichment and APIs

We understood, from relatively early on in the prototyping, that we could potentially deliver more value in terms of testable prototypes if we brought additional data into the records, and explored using machine methods to link records to each other via *entities* (people, places and organizations) that occur across many records.

Approaches

Named entity recognition

One simple approach that we have used before on previous projects is to use natural language processing tools such as [Spacy.io](#), [NLTK](#), or [BERT](#) to identify named entities in catalog record text fields, and use these to link records together via shared named entities.

This has worked well on previous projects where we have a relatively small corpus of records, and a relatively small group of possible entities because it is possible to disambiguate entities and make informed decisions about *identity* relations between entities, i.e. that is that *this* term in *this* record is the same term as *that* term in *that* record. In those types of project, the key identifiers may well be *internal* to the project rather than widely used linked data sources.

However, with the LoC data, we knew that:

1. We had a much larger and more heterogeneous collection, and little or no time for tuning named entity recognition or tying that recognition to a known set of entities.
2. We wanted to link entities to other sources of data such as Wikipedia, Geonames, shapefile data from the Census.gov, etc.

Basic naive entity recognition could not be effective, for these prototypes, without some more effective way of disambiguating entities and linking them to identifiers shared by sources of linked open data.

Entity linking

For these prototypes, instead, we needed to both identify entities in catalog record text *and* link those entities to identifiers in controlled linked data vocabularies.

We explored the possibility of building bespoke tooling to do this, and did some initial proof of concept scoping using NLP tools ([Spacy.io](#) and [BERT](#) in this case) in combination with public APIs such as Wikidata’s API to disambiguate entities. This was relatively effective.

However, it became clear, relatively quickly, that we would have better results by leveraging existing entity linking tools such as [OpenTapioca](#) or the [Radboud Entity Linker](#)⁷.

We tested several approaches, and in our initial Fat Item Page prototype (the [Iteration 1: “Skinning the LoC JSON API”](#) version, see above) we tested both the [Radboud Entity Linker](#) and [OpenTapioca](#) in our on-the-fly enrichment to evaluate how useful the data emerging from each tool was.

Our eventual prototype enrichment pipeline settled on using OpenTapioca to identify Wikidata IDs in catalog records. Note, that we sent the entire text of multiple fields in a record for entity linking as context is often important in entity disambiguation, and sending individual fields would be less likely to produce accurate results.

Entity enrichment

The output from entity linking, however, is not enough to enrich an item page (to produce the *Fat Item Page*) or to provide some of the visualizations (maps, for example) that we wanted to add to both prototypes.

However, once we have an identifier, for example from Wikidata, we can use *that* identifier to crosswalk across multiple data sources and fetch data from:

- Wikidata
- Wikipedia
- GeoNames
- US Census shapefile data

We wrote integrations for all of these, and for some additional APIs (such as the Getty Thesaurus of Geographic Names).

⁷ See <https://github.com/informagi/REL> for usage and example APIs.

Technology stack and queue based processing

Order of operations

The typical process of ingesting a collection into the prototype APIs was as follows:

1. Fetch collection records from LoC JSON API
2. Create Item records in the prototype Item API
3. Extract a list of Wikidata entities from Item record data using OpenTapioca (or other entity linking/disambiguation tool)
4. Fetch Wikidata data (for each entity) from Wikidata using the MediaWiki API
5. Fetch additional data from Wikipedia (for all entities)
6. Fetch data from US Census files (for US geographic entities)
7. Link entities together⁸
8. Link entity data to Item records
9. Index Entity data AND Item record data into a separate Search API (to drive aggregation and clustering, and linking between Items)

Signals and tasks

It made sense to keep these operations relatively decoupled from each other. In some cases an operation might fail as some third party API might be unavailable or throw an error. In other cases, one source of data might exist for some entity but another source of data may be unavailable.

Each of our data models was able to produce and respond to signals from related models and these signals could spawn one or more asynchronous tasks in our task queue for processing by multiple workers.

In this way, coupling between tasks and between data models in the underlying database application was able to flex effectively to changing data without requiring rigid coupling or the time-consuming explicit handling of edge cases.

Asynchronous task processing

⁸ The backend entity API knows, for example, that a particular geographic entity identified via a FIPS identifier is the *same* entity as a geographic entity identified by a Wikidata identifier. The API also, in some cases, knows that a particular entity is *within* another entity.

Processing a collection of a few hundred items from the LoC JSON API could, potentially, involve many thousands or even tens of thousands of operations, each of which might require an HTTP request to a public API. One item, for example, might be linked to dozens of entities, each of which might require multiple HTTP requests to fetch and populate any enriched data for that entity.

HTTP requests intrinsically lend themselves to asynchronous processing. Each request might involve hundreds of milliseconds of idle time between making the request and receiving data. Spread over many hundreds of thousands of requests this is a significant problem. Instead, it makes more sense to process these asynchronously via queues which can feed multiple *workers* each of which can be fetching tasks from the queue, fetching data, and then parsing and processing that data. While one *worker* is idle, waiting for data, the other workers are continuing to process other tasks from the queue.

Most of the time, our stack ran multiple concurrent workers reading from the task queue, although this did occasionally pose problems with the LoC JSON API (see [LoC APIs](#) below) which is rate limited. We were able to vary the number of workers throughout the project, ramping up to do more entity enrichment, for example, after we had ingested a large number of objects.

The queue broker, in this case, used PostgreSQL as the backend, and Python to define the asynchronous tasks so that it was easy to deploy and run in a local development environment. For a larger non-prototype project, it would potentially make sense to use a different enterprise class set of tools for queue handling, although the basic architecture and approach would be very similar.

Search

The final stage of our processing workflow was to index data into a dedicated single-purpose search API. This search API indexes atoms of content such as individual metadata fields, entity labels and identifiers, etc into a simple database architecture which can be heavily optimized to efficiently find results⁹ and to calculate facet counts across records.

The search service is relatively data model agnostic, because it treats each field as a discrete piece of indexable content which is normalized to a simple underlying text, date or numerical format.¹⁰ This makes it possible to change the underlying models more or less at will throughout the prototyping process, at the cost of some efficiency at ingest time.

⁹ In this case, the PostgreSQL database implemented full text indexing using precomputed search vectors, and employed B-Tree, GiST, GIN and other indexes for fast indexing of data.

¹⁰ The basic approach is similar to an [Entity-Attribute-Value](#) store, with the same advantages and disadvantages.

Enrichment and APIs: Risks and challenges

There were a number of challenges in the data storage and enrichment process.

LoC APIs

One problem was that the Library of Congress JSON API has very aggressive rate-limiting. Any request rate much higher than one request per second, or involving parallel requests (possible with our multi-worker task queues) resulted in regular [HTTP 429](#) errors. On some occasions we also noticed regular [HTTP 503](#) errors, too. This made the harvesting of collections difficult, and fragile, with collections often only partially harvested rather than harvested in their totality.

Our harvesting architecture could requeue missed items, but this often resulted in a cascade of 429 errors.

In retrospect, it would have been good to have explicitly flagged this problem with the LoC earlier in the process, and found a workaround. We would anticipate that this would not be a problem for a future in-house LoC project or a project implementing similar approaches with privileged access rather than access via the standard public API.

Inaccurate entity linking

In our informal internal reviews, we found that entity linking and entity disambiguation was generally fairly accurate. However, the amount of text provided was sometimes too low for the tools to accurately disambiguate between similar or related entities. In addition, most entity linking tools, including those used in this project, allow for a relatively wide range of accuracy in terms of matching in order to err on the side of not missing potential matches.

In this instance, we chose to only accept the highest ranked matches in each case, and discarded any of the lower ranked matches. This would certainly have resulted in potentially accurate matches being missed, but with the benefit of reducing the number of misidentifications in the data.

We would expect that for a larger scale “production” quality project a lot of potential opportunities would be there for training tools on known-good identifiers already present in the records, and for

tuning the entity linking, for example, by privileging certain specific fields within the records or by pre-processing the text.

“Looping” and idempotency

One disadvantage of a loosely coupled queue based architecture (see above) is that it is possible that multiple workers might be attempting the same operation, or even attempting an earlier (conceptually) operation in the same flow. The same operation might happen multiple times.

For example, a worker carrying out entity enrichment on a specific entity (from Wikidata) might be fetching Wikipedia data that had already been fetched by a different worker. Or a change to that entity record might trigger linking to another entity or fetching of geodata from US Census files where that operation had already happened. In the worst case scenario data might even be lost if an operation reset some underlying data model instance to an early state.

This was relatively easy to work around by elaborating slightly on the conditions that trigger certain tasks.

However, it was important to ensure that almost all task driven operations were [*idempotent*](#), that is, able to be run multiple times without any knowledge of the state of the underlying object.

Prototype 3: Tagging and Sharing

Problems and proposed solution

We know users want to collaborate with colleagues or other organizations to get work done. Also, being able to see they are having a positive impact on people is a big motivator for them. Providing a simple way for all users to contribute their knowledge to LoC's collections and sharing that beyond the LoC platform could facilitate networking, setting structures for learning and ultimately making a positive impact.

User Insight 1: Collaboration + Impact

- How can users collaborate with others?
- Share their contributions?
- See the positive impact their work has on others?
- How can users get help from other users, who might have a different level of expertise or knowledge?
- How can users share their work with others?

Proposal

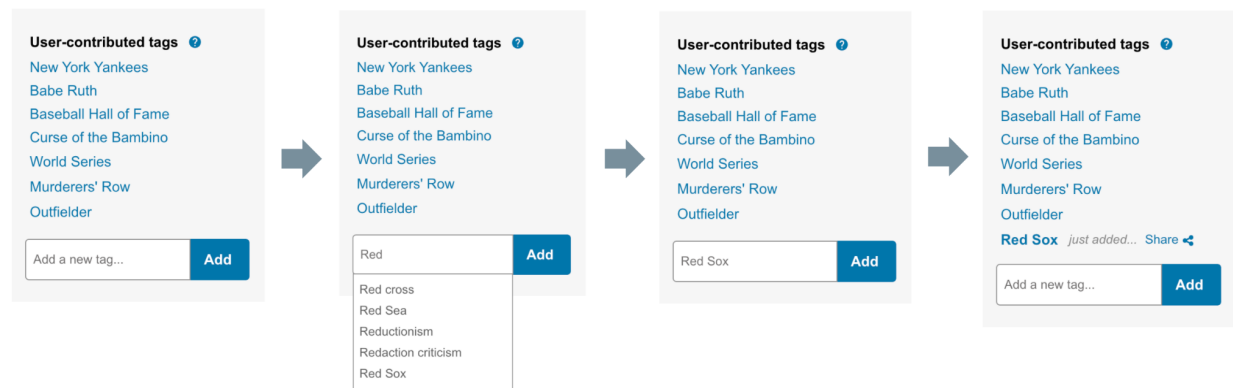
This is really two linked proposals:

1. Catalog pages and search result clusters could have user contributed tags. These could autocomplete against LCSH or similar vocabularies, or just emerge organically from users.
2. Sharing should be made better by really concentrating on what is made available, for example on social media platforms, when a user shares a LoC resource, such as:
 - a. a tag they have made,
 - b. a search they have performed, or
 - c. an item they have found

Process: Wireframing

Unlike the other two prototypes, we focused here on producing a number of [interactive wireframes](#) to help test these ideas with end users. These wireframes were sufficient to validate the concepts with users and allowed us to spend more development effort on the other two data-driven prototypes.

Tagging



Rather than go down some deep sharing/contributing/crowdsourcing platform approach, we allow simple user generated tags to organically generate clusters of documents, and then we show these on the catalog pages, but we make sure the provenance of them is clear. That these are user contributed.

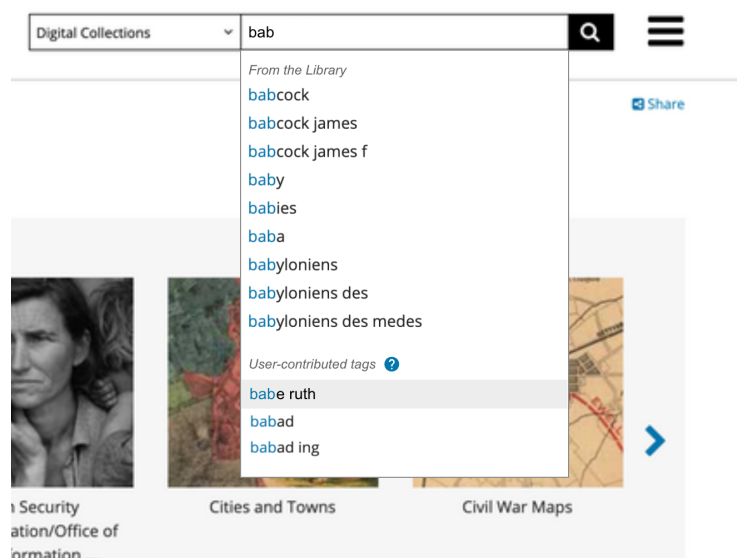
Users get very immediate feedback from adding tags, as they can clearly see their work—the tagging and identification of objects—being made available to others.

The key thing is that the features are very limited, but we make it very easy to do, and we provide good aggregation/grouping and onward journeys from each tag, e.g. by showing related tags, or showing where other people have tagged similar documents, or used similar tag terms.

There is then a clear development path to move from this—simple tagging—to user generated collections, and user curated content that goes beyond simple tags. Tagging is a good way to start the process and work through issues like UX, approval and review, etc.

By design this widget displays a maximum of 10 items at a time. We have documented other scenarios including when that number is exceeded in the [wireframes](#).

We also explored how these user contributed tags can enhance the search experience provided that these terms are clearly highlighted as user contribution data:



Sharing

The Library already makes use of tools such as Open Graph metadata embedded on pages to facilitate sharing on social media and on other platforms. Rather than make the Library into a social media or sharing platform—which is going to be extremely hard to do when competing with massive social media platforms—we should facilitate better sharing of LoC material on social media, and in the external tools that users actually use.

We would propose to enhance this existing feature to create a richer display of data, using the same Open Graph tooling already used by the Library, with thorough user testing of what is most compelling and useful to them when they encounter LoC material on social media.

We would also propose that the Library consider how they might use formats such as RIS or other formats supported by reference managers and research tools—tools such as Evernote, Mendeley, Zotero, etc.—to make it easy for users, including expert users such as data journalists, to store, reuse, and cite material they find on the Library site.

We explored a consistent presentation for whether a user is sharing a catalog item page, a newly user contributed tag, a collection page or search results. These examples can be seen [here](#).

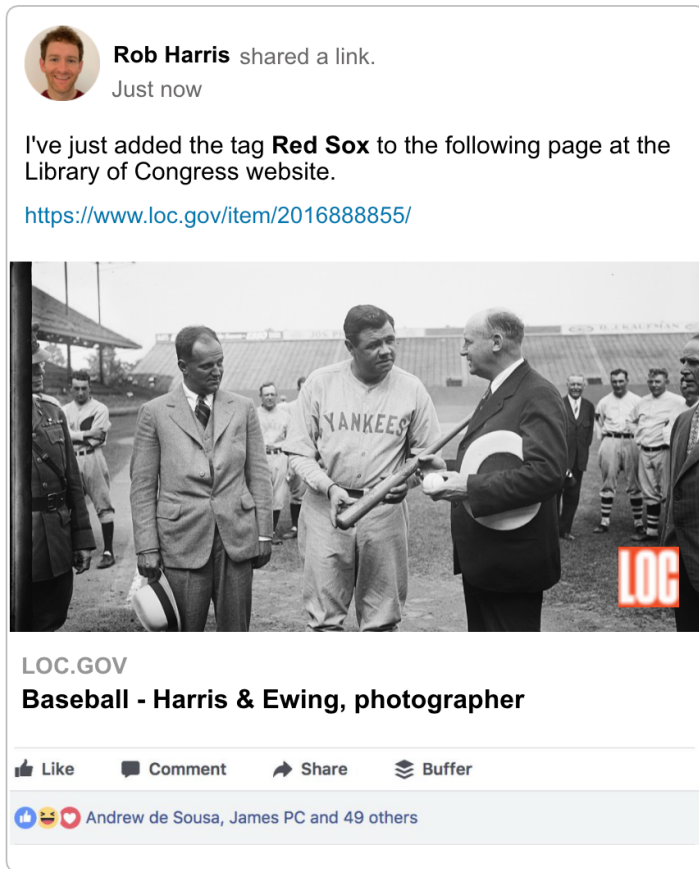


Figure. Catalog item page Open Graph sharing.

User insights

These interactive wireframes focused mainly on validating the idea rather than identifying usability problems hence the focus solely on user insights:

User-generated tags upfront

The wireframes we used to test this idea placed this tagging widget below existing filters and metadata sections of various page types in the LoC site.

People suggested that this widget should have more prominence in the page hierarchy. Otherwise it would get buried down a long list of LoC metadata information in their view.

A complement to fat item pages

Some users asked why we didn't include this feature in the fat item page prototypes. They could see the benefit of these two concepts working in tandem.

Quick feedback

The ability for added tags to be immediately available for the creator of the tag and others to use in order to find content was important to users. They believe it could make people feel more involved with the Library.

Great for learning concepts from users

One user thought how beneficial this could be for the Library when it comes to understanding how members of the public might classify content.

Trust and value

The same user asked whether LoC would need to periodically move these user generated tags to subject levels.

Managing volume

One user who in the past had to filter lots of user generated content was cautious on the large scale implementation of this. Another user suggested to only 'promote' those tags that were added by a number of people as a sort of validation.

Why stop at tagging?

A user suggested why leave it at tagging? Why not let people add their own stories if they were part of it. *'My mum knew X...'*. This user wondered how to deal with fake info though but thought most people would like that feature.

Rewarding users

A user felt that those contributing tags to collections and items might appreciate some recognition like a public profile within the LoC website.

Challenges and lessons learned

Data Retrieval and Enrichment

What went well

Our basic approach to fetching data and enriching data worked well. The core model with:

- Loosely coupled data model instances and processes
- Asynchronous tasks triggered by signals on data model instances
- Multiple workers pulling tasks from a task queue

Was effective, and scaled well to large-ish workloads.

On the enrichment side, bringing together data from multiple 3rd party sources including Wikidata, Wikipedia, GeoNames, and the US Census was relatively straightforward to produce effective results given the use of shared identifiers.

The LoC JSON API data model was easy to work with, and the search API was also easy to work with when implementing the [Iteration 1: "Skinning the LoC JSON API"](#) prototype.

Challenges

The LoC JSON API was a challenge for bulk data retrieval. Aggressive rate limiting meant that it was never really possible to harvest records reliably at scale or at the rate that would have been possible using even with a relatively modest backend setup on our cloud-hosted version of the prototype platform. HTTP 429 or 503 errors were common, and we were forced to requeue jobs continually or introduce artificial "sleep" periods after every HTTP request to avoid being throttled. This was particularly problematic after we ended core development on the prototypes and began to reharvest data into a new instance of the platform for user testing. In retrospect it would have been better to have addressed this with the LoC at an earlier stage.

Data quality, accuracy and provenance

What went well

We did not do a detailed manual audit for the accuracy of entity linking across the corpus of records we ingested from the LoC API. However, as a general rule, our impression was that the accuracy was more than adequate for a prototype of this type. Our use of entity linking erred on the side of caution—only accepting the highest ranked potential match, in each case—and this helped.

We would expect that in a future project that made use of similar or related methods more attention would need to be paid to data accuracy. However, we are confident that there are improvements to be made in advance of any human-in-the-loop refinement or curation of machine-generated data. Our use of machine enrichment methods was time-constrained during this project, and we were more concerned with user experience than with data, except to the extent that data could be used to facilitate the creation and evaluation of the prototypes.

The current state of the art for entity recognition, entity disambiguation, and entity linking is relatively fast moving, and as a team we were aware of multiple potential good candidates—some with more flexibility and potential for extension or customization—that we didn't test on LoC material. There is a lot of potential for improvements in accuracy and comprehensiveness.

Challenges

Important concerns raised by LoC staff during testing included:

Accuracy of the data

Standards for data accuracy are, naturally, much higher when records are being retrieved from an institution of record like the Library of Congress. Cataloging is a skilled and time-consuming process, and machine methods are not likely to replicate this level of care and accuracy any time soon.

Data provenance

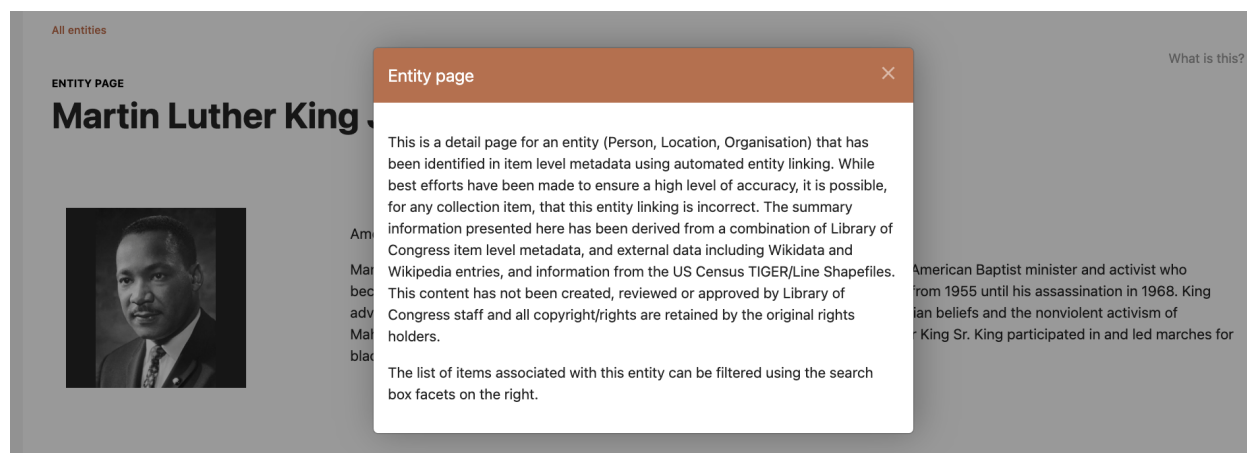
Our initial prototypes did not make clear on the page what the source of the data was, nor did they distinguish between data created by the Library and data created via machine methods.

Contested, sensitive or pejorative terms

Using machine methods runs the risk of *surfacing* existing terms in the catalog records that may be contested, politically sensitive, racist or culturally sensitive, sexist, or reflective of some historic

bias or inaccuracy. These methods also run the risk of introducing *new* contested, sensitive or otherwise incorrect or inappropriate terms.

In the current prototype we addressed all of these concerns through an information box and disclaimer available via the “What is this?” link on many of the page elements.



What is this? On an entity page

While these pop-up information boxes do a relatively good job of conveying *how* the information was created (process) and *where* the information came from (provenance) they do not address accuracy or issues around pejorative, contested or sensitive terms, except as a disclaimer to advertise that LoC staff were not involved.

Future work in this area would need to address accuracy, provenance and contested or sensitive terms in a more robust way (see below).

User Research

Challenges

User recruitment for both rounds of user research was difficult, and ideally we would have had access to a larger and more diverse pool of users for research and user testing. In future projects, it would be beneficial to consider involving user recruitment specialist agencies early to mitigate any risks and ensure the project doesn't get delayed because of lack of users to test with.

Discovery and Ideation

What went well

The overall approach, from the Discovery and User Research phases, working closely with LoC staff to inform the Ideation and Sketching sessions all went very well. When work began on the prototypes there was a good shared sense, within the development team, of what the priorities were for LoC Users (via the User Insights) and where the most promising opportunities lay (via the How Might We sessions).

Most of the ideas that emerged and were prioritized during sketching were prototyped at one level or another, either as standalone frontend prototypes (via CodeSandbox and Semiotic) or integrated into one or other of the dynamic prototypes fed by APIs on the [main prototype site](#).

Challenges

Some of the prototypes that we would have liked to have taken into the final prototype site, such as the [facet co-occurrence matrix](#), were left as standalone prototypes. This was largely due to the limited time available, and a conscious choice to try to integrate the prototypes into the search and entity pages, rather than leave them as unconnected pages..

Development approach

What went well

The broad approach across the frontend and backend applications went well. Using React for the frontend and standalone tools for hosting and development like CodeSandbox was a productive way of maintaining momentum on the front and backend sides of the project simultaneously and enabled backend developers (using frameworks like Semiotic) to quickly test how API changes integrated with approaches to visualization.

Using a relational database with REST APIs was also productive in terms of relatively rapid prototyping of API changes and the modeling and serialization of data.

Challenges

Some of our prototyping on the frontend was not carried through to the final prototypes. This is always a potential feature of a fast moving Agile “alpha” project of this type.

Using a relational database was very efficient for development, but only adequate for performance in the final prototypes. If a similar project was being done with an emphasis on run-time query performance and with a “production” mindset, it would potentially make sense to:

1. make more use of preprocessed data with the data model(s) reified in a fast search solution like Elasticsearch or Solr, or
2. to “bake in” the data model into search (within a relational database) in a way that privileges performance over flexibility (whereas the prototyping process deliberately chose to do the reverse), or
3. Adopt some hybrid of these and other approaches to storing and indexing large quantities of data.

Lessons Learned

When it comes to addressing some of the challenges or opportunities identified in the preceding sections, there are a number of potential avenues that might be worth further exploration.

Accuracy, provenance and sensitive themes

In terms of data accuracy, we spent very little time working on data accuracy as part of the prototype. Accuracy was “good enough” for the prototyping process and for user testing.

There are obvious opportunities in terms of:

- Evaluating a wider range of existing entity linking and entity disambiguation tools/libraries
- Training entity disambiguation tools on LoC data with known good identifiers
- Indexing data *within* objects, such as text for documents available with transcripts or which are amenable to OCR, as well as data *about* objects
- Preprocessing text or data records to ensure only relevant fields and relevant textual data are fed into the entity linking and disambiguation
- Iterative improvement of models and processes via “human-in-the-loop” style processing in which scoring and review of a subset of the machine generated data is used to improve the overall corpus of machine generated data

When it comes to risks associated with machine enrichment either surfacing or creating pejorative, culturally sensitive or otherwise problematic terms, there is almost certainly no substitute for some level of data curation and review.

However, this need not necessarily involve record by record review of data outputs. Instead, review of the themes identified by machine methods could be carried out globally—by LoC staff

and subject experts—at the level of themes that span multiple items, and the outputs of this review fed back into the tagging process.

N.B. It is important to note here that the entities identified in the prototyping process were restricted to people, places and organizations. These are intrinsically less likely to be problematic than other classes that might be identified, such as subject terms or conceptual themes. While there are still risks—such as, for example, when an offensive term for a group or organization is used rather than the preferred term¹¹—there are far fewer risks with concrete entities of this type than there are with more abstract terms/entities.

It may also be worth extending and revising the signposting in the UI that presents users with caveats about data curation and machine methods to include a feedback mechanism to solicit input from users about problematic terms.

Performance and scalability

The broad approach taken on the project: loose coupling of data, asynchronous task queues with multiple workers for data processing, etc. would all make sense on any future project, and are standard approaches to working with data at scale. However, the prototyping was obviously limited by time and budget, and was primarily not concerned with overall performance, scalability or robustness.

Obvious areas for improvement might include:

- Data storage and indexing using Elasticsearch or Solr (or another enterprise scale search solution)
- Precomputing computationally expensive data such as, for example, clustering on large collections or other analytic tasks not suitable for run-time computation
- Highly parallelized enterprise level stream/batch processing for data extraction, transformation and data load

Usability of the prototypes

There were some minor usability areas discovered during the testing but that was expected at this stage of the development. These issues, as well as suggestions on how to address them, are documented in this report - see Usability notes sections for prototypes 1 and 2.

¹¹ Digirati have worked on projects with Native American treaty records from the US National Archives, for example, where some of the formal labels for tribes *in the source documents* use terms that are now considered pejorative, offensive or historically inaccurate by those tribes.

Also worth noting is that there appears to be some differences about what LoC might consider relevant data about an item and the end users' view on that. Library specific vocabulary for labels might make that gap bigger. This is also documented in more detail in the Usability notes sections.

Conclusions and potential next steps

Conclusions

The user types we targeted are looking for a user experience that is more tailored to them as a user, and less tightly tied to the vocabularies and conceptual schemes that are familiar to librarians and information experts. This was not extensively explored during the prototyping process and is a promising area for further prototyping and user testing.

Creatives and data journalists are looking for a search experience that gets them to the results they want more quickly and with less trial and error or cross-referencing with external tools like Google. Other users, however, such as students and educators, are interested in a more exploratory interface, and the site needs to support both classes of user. Users generally are looking to see more information at their fingertips with less clicking, and rich information is welcome, which validates some of the approaches tested in the Fat Item Page prototype, in particular.

Visualizations, as found in the **Cluster and Aggregate Search** prototype, were promising but need further refinement and testing with users to identify which types of visualization, and which labels or interactions, are most helpful. Users commented on the tension between search results and additional explorative artifacts so we recognise that iterating more often on this area would have been helpful to get the balance between exploration and search right. But, as highlighted in the Challenges and lessons learned section, we were constrained by limited access to users that fit the targeted audience. We'd recommend keeping in touch with these audiences - some people interviewed expressed interest to continue to help - and build a pool of users to consult periodically.

The **Fat Item Page** prototype was the best received of the three prototypes. Users saw the benefits this type of rich but well-organized resource could bring to them and others. Educators in particular were excited about the possibilities going as far as suggesting pages like this could be used as 'national teaching instruments'. Entity extraction provides useful information for some users, especially educators, but more work is needed to refine the labels and language used to

convey to users what these are and how they might be useful. The sections on user insights and usability notes on this prototype can shed more light on what to address going forward on this. The **tagging and sharing** component was seen as an additional feature that a fat item page should have.

Potential next steps

If the Library wishes to pursue some of the concepts and ideas derived from the user research and tested through the prototypes—and based on the findings from the user research and user testing this would be our recommendation—there are a number of initial concrete next steps that could be taken to develop these ideas further.

1. Enrichment

As per the previous section, there are many potential improvements that can be made to improve the efficiency, accuracy and comprehensiveness of the machine generated data. This could include: evaluating a wider range of tools; training tools on LoC data; more data curation and “humans in the loop” feedback; and better signposting in the UI with clear indications of the provenance, and likely accuracy of the data, with opportunities for users to provide feedback or flag issues.

2. Performance

These prototypes were primarily intended to evaluate approaches to user experience and to approaches to creating, storing and presenting data to service those user experiences. While performance had to be adequate for user testing, little optimization was done for performance at an enterprise-level scale, either in terms of data processing, or in terms of run-time search performance. If elements of these prototypes, or the approaches they explored were adopted by LoC, then things could be done differently. There are many tried, tested and well understood approaches that could be adopted in future projects that make use of some of the same ideas to: improve run-time query performance; pre-compute computationally intensive data; improve the efficiency of data storage and retrieval, and ramp up workflows for data processing.

3. Visualization

Some of our visualization experiments did not make it into the final prototypes. There are a number of promising avenues that could be explored here. The idea that collections and known aggregations of data could be provided with a range of “infographics” could be explored in more detail, and more use could be made of computationally intensive data transformation to create some visualizations in advance, rather than relying on run-time queries in every case.

4. Personalization

Users expect to see content that matters to them. They told us that. They mentioned how important it is for the content, its organization and labeling to work in tandem to deliver truly personalized results. So this could be explored further, perhaps looking at how taxonomy-driven personalization could be implemented as the foundation of a personalization strategy, based on the LoC Catalog and utilizing the enrichment techniques explored in this project.

5. Discovery vs. Search

Both approaches are intrinsic to content-rich online experiences. These prototypes have expanded the possibilities of exploring relationships in the LoC catalog and beyond and that excited users. However, getting the balance right between exploration (or discovery) and search should be an ongoing task for LoC. While people like students and researchers can benefit from discovering new items and collections in the catalog, others with time constraints need a powerful and accurate search capability to get the content they need fast. This is also an area that tightly links with personalization.

6. Research wider audiences and continue learning

The ideas developed in these prototypes can also have an impact on the experience of users we have not considered in this project. We would recommend the Labs team and LoC as a whole to continue exploring these concepts further, involving wider user groups and keep learning on the potential these prototypes have. Other user groups that were aligned to the project goals were considered early in the discovery phase (see diagram in page 11) but also wider audiences can be studied from revisiting the inverted pyramid of user engagement model we introduced at the beginning of the project (see diagram in page 10).

Appendix 1: Prototype

“pitches”

Cluster and aggregate search

For

All of the core user personas

- Educators
- Activists / community leaders
- Data journalists
- Undergraduates / creative arts students

Their problem(s)

This particular set of problems impacts on all of the core user types. Carrying out searches and understanding the results of those searches is a fundamental requirement for almost every user in their journey to discover and use the Library’s collections. This requirement is particularly difficult to deal with when users are either inexperienced at using and understanding library catalogs, or when users are under time pressure.

User Insight 3: Search and Ask a Librarian

- Users struggle to make sense of search results
- Users struggle to find what they are looking for
- User struggle to understand what the Library holds and does not
- Users fail to quickly determine whether there is an answer to their question
- Users go to Ask a Librarian when search could have answered their questions, but they were unable to find what they wanted

User Insight 4: Today Problem

- Users do not have time to spend understanding complex or difficult to comprehend resources
- Collections are not optimally organized for time-pressured users

- Users want to quickly and effectively access content

Proposal

The data in the catalog can be understood as a rich multidimensional nested graph of information, so we could explore multidimensional ways to represent this rich graph-like data, rather than as a unidimensional flat list.

Instead of presenting the search results as a single ordered list with facets, present the information:

4. Clustered hierarchically into groups
5. With useful summary information about each group in aggregate
6. Employing techniques from data visualization to make the results easier to comprehend at a glance

So that a user can browse/navigate to the group of results they want and then from there to more specific clusters and individual records.

Currently, search results are represented as:

- A. A single flat list
- B. Information about that list is either:
 - a. A count
 - b. Facets with counts per facet
- C. The user journey goes directly from list to record
- D. Filtering applies to the whole list at once

Users have little sense of what might lie beyond the first page of results, and facets are powerful as a tool for filtering information, but not necessarily the best way of presenting information to a user in order to help them make decisions and continue their search journey.

Instead, we could present those results as:

- A. Clusters of results, with more than one cluster visible at a time. Clusters could be
 - a. Natural clusters derived from shared common metadata fields, potentially with many fields being used to cluster these results, rather than single value / single field, as per facets

- b. Generated from resource content, such as OCR data or image recognition, so that clustering goes beyond the catalog data
- B. With summary information about the clusters in aggregate, that goes beyond a single “count” value, and tells the user, at a glance, something useful about what they might encounter if they drill further down into a specific cluster
- C. Hierarchically or “topographically” through the graph of information, with users navigating into clusters, and potentially being presented with further clusters, so that they can navigate/browse towards the specific resources they want. Navigating, not necessarily through strict collection hierarchies, but rather through groups of related concepts, or from the general to the more specific.
- D. Using visualization and summarization to convey information to the user at a glance

How can it help the Library of Congress?

This proposed prototype addresses a number of the How Might We statements developed with LoC in the earlier ideation and prioritization session

- HMW get rid of search result pages altogether and find another way for users to explore/find relevant info?
- HMW extend search beyond the catalog record paradigm?
- HMW provide clear paths to the refinement of search?
- HWM improve the relevance of search results?
- HMW demonstrate what is often used in relation to the user's search/question
- HMW remember what it is like to be a beginner?
- LoC as Primary Resource (because it's a quick and easy to use source of information)
- (not specifically chosen in prioritization but relevant) HWM reduce the amount of reading a user has to do?
- (not specifically chosen in prioritization but relevant) HWM summarize and aggregate information for speed of comprehension

How might it work?

The library already has APIs for fetching catalog data, documented by the Library Labs as part of LC for Robots, and has a JSON API for search. We could build upon and enhance this using algorithms and machine methods to cluster data, or to enhance the clusters that are explicit or implicit in the user's search terms to transform flat lists of results into clusters and hierarchies. We could use machine methods to extract data from content, where available, such as by

generating keywords or topic clusters from OCR data, so that the results extend beyond the catalog record and into the object itself.

We could use data visualization techniques and good UX design to try to convey summary information to users so that they can quickly understand what they are looking at, at a glance, and to aid them in choices they make on their search journey so that they can more quickly get to the objects they want.

We could retain the ability to switch back to a list view, either at a certain point in the user's journey, or at any time, via some UI toggle, so that the user is not locked into only one search paradigm.

Resources and supplementary material

<https://opus.lib.uts.edu.au/handle/10453/135763>

<https://www.vam.ac.uk/blog/digital/visualising-library-catalogues>

<https://searchisover.org/papers/albrecht.pdf>

Visualization

N.B. presented here as simple examples, and not intended to represent what we'd build.

<https://observablehq.com/@d3/zoomable-circle-packing>

<https://observablehq.com/@d3/zoomable-icicle>

<https://observablehq.com/@d3/tidy-tree?collection=@d3/d3-hierarchy>

Other issues and considerations

What is the appropriate scale to work at, when prototyping? Is there a balance between comprehensiveness and tractability (as a practical problem during development).

Fat Item Pages

For

All of the core user personas

- Educators
- Activists / community leaders
- Data journalists
- Undergraduates / creative arts students

Their problem

Users searching for content ultimately end up on a catalog or item page. Item pages can be quite rich sources of information but they may not always be easily comprehended by a user. Nor do they always contain all of the information that can potentially be generated about that item, or links to all of the associated sources of information or related items.

User Insight 3: Search and Ask a Librarian

- Users struggle to make sense of search results
- Users struggle to find what they are looking for
- User struggle to understand what the Library holds and does not
- Users fail to quickly determine whether there is an answer to their question
- Users go to Ask a Librarian when search could have answered their questions, but they were unable to find what they wanted

User Insight 4: Today Problem

- Users do not have time to spend understanding complex or difficult to comprehend resources
- Users want to quickly and effectively access content

User Insight 2: Trust

- Users want trustworthy data.
- Users want to rely on LoC to help make their information-driven decisions and arguments.

Proposal

Item pages should not just present the Marc record and a simple embedded viewer back to the user with some limited linking to related records via subject terms.

Item pages could be “fat” pages rich in sources of data to provide the user with the most information to help solve their information-driven needs, and to provide them with clear onward journeys to likely other sources of data if this item page does not provide them with what they want.

Item pages could bring data in from other sources, such as:

- LoC collection homepage information
- LoC linked data/ authority sources, such as LCNAF
- Information from LoC research guides
- Rich crosslinking (generous interfaces) to related records within the catalog, potentially generated via machine methods such as named entity recognition, and named entity linking
- (On the ecosystem / trust space) we might link out to other records in related US governmental sources like archive.gov (which has a JSON API, so this could be done)
- Potential 3rd party sources such as Wikidata or wikipedia
- Presenting summary information and aggregation information about related records

Even if this results in repetition across catalog pages, the beginner user, or the person beset by the “Today problem” isn’t likely to look at many item pages, and will appreciate the rich information being presented to them.

We could also have some UI affordances to turn up or down the amount of additional data.

How can it help the Library of Congress?

This proposed prototype addresses a number of the How Might We statements developed with LoC in the earlier ideation and prioritization session

- How do we remember what it is like to be a beginner?
- HMW extend search beyond the catalog record paradigm?
- HMW enable collaboration with other organizations both in the private and public sector to raise LoC profile as a source of trustable information?
- HMW facilitate comparison to information in other collections or other organizations
- (not singled out but part of the Today problem theme) HMW better integrate LoC resources in broader web search (knowledge graph, etc)
- (not singled out but part of the Today problem theme) HMW make LoC a first-stop when looking for information in a hurry?

How might it work?

The library already has APIs for fetching catalog data, documented by the Library Labs as part of [LC for Robots](#), and has a JSON API for search. We can bring these records in and use:

- Identifiers already present in the records
- Relationships already present in the records
- New identifiers and new relationships that we add using named entity recognition and named entity linking

To bring in data from other sources within the LoC datasphere, and also from other sources.

We can explore using different methods of presenting the data, for example by using summarization, or visualization to return rich data in quick, easy to comprehend ways that we can embed within the page.

N.B. This proposal is as much a data presentation, visualization and design/UX task as it is a data enrichment and data linking task. Methods for enriching data are well understood and Digirati have experience in using these, but the success of the prototype rests on being able to present this data to users in ways that work for them and address their data needs.

Tagging and sharing

Who is it for

All of the core user personas

- Educators
- Activists / community leaders
- Data journalists
- Undergraduates / creative arts students

Their problem(s)

We know users want to collaborate with colleagues or other organizations to get work done. Also, being able to see they are having a positive impact on people is a big motivator for them.

Providing a simple way for all users to contribute their knowledge to LoC's collections and sharing that beyond the LoC platform could facilitate networking, setting structures for learning and ultimately making a positive impact.

User Insight 1: Collaboration + Impact

- How can users collaborate with others?
- Share their contributions?
- See the positive impact their work has on others?
- How can users get help from other users, who might have a different level of expertise or knowledge?
- How can users share their work with others?

Proposal

This is really two linked proposals:

1. Catalog pages and search result clusters could have user contributed tags. These could autocomplete against LCSH or similar vocabularies, or just emerge organically from users.
2. Sharing should be made better by really concentrating on what is made available, for example on social media platforms, when a user shares a LoC resource, such as:
 - a. a tag they have made,
 - b. a search they have performed, or
 - c. an item they have found

We would propose, in this instance, that this proposal be primarily about wireframes and design/UX and user testing work, rather than a fully working prototype, although some limited version of tagging and sharing could be implemented as a proof of concept on “Fat Item Pages”.

How can it help the Library of Congress?

- HMW we connect groups in technical and physical spaces? (In this case, by making it easier for users to create collections and tag objects, or to take LoC objects into their preferred social media or research platform)
- HMW: Allow a user to share curated collections of material from the LoC?
- HMW provide a digital space in which collaboration can happen?

How might it work?

Tagging

Rather than go down some deep sharing/contributing/crowdsourcing platform approach, we allow simple user generated tags to organically generate clusters of documents, and then we show these on the catalog pages, but we make sure the provenance of them is clear. That these are user contributed.

Users get very immediate feedback from adding tags, as they can clearly see their work—the tagging and identification of objects—being made available to others.

The key thing is that the features are very limited, but we make it super easy to do, and we provide good aggregation/grouping and onward journeys from each tag, e.g. by showing related tags, or showing where other people have tagged similar documents, or used similar tag terms.

There is then a clear development path to move from this—simple tagging—to user generated collections, and user curated content that goes beyond simple tags. Tagging is a good way to start the process and work through issues like UX, approval and review, etc.

Sharing

The Library already makes use of tools such as Open Graph metadata embedded on pages to facilitate sharing on social media and on other platforms. Rather than make the Library into a social media or sharing platform—which is going to be extremely hard to do when competing with massive social media platforms—we should facilitate better sharing of LoC material on social media, and in the external tools that users actually use.

We would propose to enhance this existing feature to create a richer display of data, using the same Open Graph tooling already used by the Library, with thorough user testing of what is most compelling and useful to them when they encounter LoC material on social media.

We would also propose that the Library consider how they might use formats such as RIS or other formats supported by reference managers and research tools—tools such as Evernote, Mendeley, Zotero, etc.— to make it easy for users, including expert users such as data journalists, to store, reuse, and cite material they find on the Library site.

Resources and supplementary material

<https://flickr.com/> (tagging)

<https://makingscience.royalsociety.org/s/rs/page/welcome> (LCSH based tagging)

[https://en.wikipedia.org/wiki/RIS_\(file_format\)](https://en.wikipedia.org/wiki/RIS_(file_format))

<https://ogp.me/> (LoC already makes extensive use of this, the proposal is about improving and enhancing through user testing)

<https://www.rijksmuseum.nl/en/rijksstudio>

Appendix 2: User research and usability test methodology

Sessions

The testing sessions were scheduled between January 24 and February 4, 2022. They were conducted over Zoom using the web-based prototypes we developed at Digirati. The video and audio of each test were recorded while notes were also taken during the sessions.

Participants

Six people took part in these sessions:

- A self-employed, multi-talented artist that has been working on a number of roles: investigative journalist, musician, author, technical writer, project manager, data journalist, poet, music producer/mixer and content creator.
- An artist and musician that produces music videos containing video and photos about black history and voting.
- A former HS Teacher, currently a digital learning design specialist who engages with library-like materials to assess them in terms of their readiness for classroom use.
- A head cataloger for PBS's educational platform (PBS LearningMedia), pivoting PBS generated materials (videos, digital documents, images, articles) to make them classroom ready.
- A research librarian at The Catholic University of America that specializes in Interlibrary Loan and Document Delivery, both for the library's patrons and requests from other libraries.
- A research Librarian at the Pratt Institute teaching Reference to Library and Information Science students.

Approach

Showing users an early-stage prototype is partially usability testing and partially user research. Our approach in this instance looked to get both feedback on the actual designs and to learn more about the users' mental model through deeper questioning. We wanted to expand where

possible on what we learnt during the first round of user research which informed the direction of these prototypes in the first place.

The conversation pattern we looked to establish can be summarized into three stages:

1. **Establish frame of mind** - Participants need to look at the prototype from the right perspective. We ask them a question simply to set up the next one, to make sure they're approaching from the right mindset. By asking these questions — even ones we know the answer to — we've primed the participant for the next series of tasks.
2. **Ask about expectations** - With the participant primed, we ask them to describe what they expect to see in the product itself. This applies to a variety of situations. At the smallest level, it might be that they've pointed to a link or button they want to click. More broadly, we can ask about their expectations for the whole application. It's the priming in the first step that lets us do this.
3. **Reveal and ask about alignment** - One cardinal rule of user research is to avoid yes-no questions, like "Does this meet your expectations?". Instead, we try something like this: "How is this different from what you expected?" By asking how it's different, we encourage them to contrast what they see with what they imagined.

Creating these tests demands crafting a script that surfaces both design issues and user insights. The script can be read in the section below. It provides the initial structure but new opportunities for exploration arise during the interview itself. As participants talk, we actively listen and use the things they say as a cue to dig deeper.

Test script

< Welcome participant >

Hello (*participant name*), thank you for joining us today. My name is Diego and I'm a user researcher at Digirati, working with the Library of Congress Labs team – I'm going to walk you through the session today.

Before we'll start I have some information for you.

So why are we here? We're currently testing a number of prototypes and concepts to learn as much as we can about the way people might use them, but also understand more about your own needs in the area we're exploring.

The session will take about an hour.

Please remember, during the next hour, that we're testing early-stage prototypes, not you. Don't worry at all about mistakes or misunderstandings. If they happen, it's the fault of the prototype. And finding about these issues gets us closer to creating a great product.

During the whole test, please try to think out loud. Share anything that's in your head. Tell me what you are looking at on the screen, what your thoughts are, what you like and what don't you like, et cetera. We want to learn about your honest reactions to these prototypes.

If you have any doubts during the test, feel free to ask me any questions. However, please remember that I might not be able to answer them during the test. We want to learn about your experience with these prototypes, so I can't guide you. I'll be glad to answer all your questions after the test.

Do you have any questions so far?

Thank you for sending us in advance your consent for this interview to take place. With your permission, I'm now going to **start recording** the session. We'll record your voice, face and actions on the screen – that will give us the full picture of the session for further analysis. We're not going to share these recordings outside of the team directly working on the prototypes.

Also, I hope you don't mind, but I'm going to take notes. It helps me with analyzing the session later.

So thanks for sticking by and listening to the formalities. Let's get to do this.

< Intro questions >

- Tell me about what you do?
- What role does research or library-like materials (for example books, articles, photos, websites, videos, audio, or other documents) play in what you do?
- How familiar are you with understanding library catalogs and collections? Do you use the LoC website for example? And if so, how?

< Prototype 1: Cluster and aggregate search >

< 1. Establish frame of mind >

- Tell me about the last time you tried to search for information online for the work that you do? What sort of things were you looking for?

- Are you generally happy with the results you might find in Google or do you go one step further and use the search functionality from the website you landed on?
- Tell me more about this: what's the hardest part about searching for content using the built-in search of a website?

Thanks. I will now share with you a prototype that looks into this area. Could you visit this link and share your screen with me please:

<https://loc-testing.ch.digtest.co.uk/site/new-search?q=luther%20king>

Scenario 1

Imagine you are doing research on Martin Luther King on the LoC website. In this example you already searched for 'Luther King'

< 2. Ask about expectations >

- Walk me through the page and please tell me what would you expect to see here?
- What do you find sometimes missing on search results pages that would help with your work?
- Can you go to the visualizations at the top and tell me about your impressions on these please?

< 3. Reveal and ask about alignment >

- What do you find most useful about this page?
- What's the hardest part about using this page?
- How is this different from what you expected?

< Prototype 2: Fat item pages >

< 1. Establish frame of mind >

- Now, think about a time when you came across an online resource that provided a huge amount of information, how do you deal with it? What sort of things do you look for?
- What was the hardest part about making sense of all that information?

Thanks. I will now show you a prototype that looks into this.

Scenario 2

Let's continue from where we left before. You're still doing your research on Martin Luther King and land on this page about a photograph of the Lorraine Motel in Memphis, Tennessee. You are particularly looking to learn more about the location and people associated with this picture.

<https://loc.ch.digtest.co.uk/site/items/2010630679>

< 2. Ask about expectations >

- Walk me through the page and please tell me what would you expect to see here?
- This is obviously a rich page with lots of information: what are your thoughts about it?
- Can you now go click on any of the links in the subject section and tell me about your impressions on these please?

< 3. Reveal and ask about alignment >

- What do you find most useful about this page?
- What's the hardest part about using this page?
- How is this different from what you expected?

< Prototype 3: Tagging and sharing >

< 1. Establish frame of mind >

- Can you talk me through about the last time you have contributed your knowledge to an online resource? How was that experience for you?
- What ways do you employ to collaborate with others online?

Thanks. I will now show you a prototype that looks into this.

<https://preview.uxpin.com/f8a31cf3f983a8a026beb656e2ae282dd5df5a84#/pages/141556204/simulate/no-panels?mode=i>

Scenario 3

You are exploring LoC's digital collection of old baseball cards. You heard that members of the public can contribute their knowledge to the collection but you are intrigued about how that's done.

< 2. Ask about expectations >

- Explore the page and please tell me: what would you expect to see here that would allow you to do just that - to contribute your knowledge to the collection?

OK, great. You want to tag this collection with the words 'Red Sox. Please try to do that and describe the task.

- What do you think you'll see when you click the Add button?

Please click it and describe what you see.

< 3. Reveal and ask about alignment >

- What do you find most useful about this functionality?
- What's the hardest part about using this functionality?

< Thank participant >

Thank you, that was a very helpful session. Your feedback is of extreme help for us and the work we are doing.

Do you have any questions for us we can help with?

Thanks again for your time and have a good day.